

MC PUE 1000

CDRL #004

2

SPARTA, INC.

AD-A210 248

DDN Trusted Guard Gateway Phase 2 Report

Trusted Guard Gateway (TGG) Technology Assessment

DTIC
SELECTED
JUL 11 1989
S D

February 2, 1989

Contract No. DCA100-87-C-0095

Prepared For:

Defense Communications Engineering Center
Defense Communications Agency
Code R640, ATTN: COR
1860 Wiehle Avenue
Reston, VA 22090-5500

SPARTA, Inc.
7926 Jones Branch Drive
Suite 1070
McLean, VA 22102
(703) 448-0210

DISTRIBUTION STATEMENT A
Approved for public release
Distribution Unlimited

REPORT DOCUMENTATION PAGE

1. REPORT SECURITY CLASSIFICATION UNCLASSIFIED		1b. RESTRICTIVE MARKINGS N/A	
2b. SECURITY CLASSIFICATION AUTHORITY		3. DISTRIBUTION/AVAILABILITY OF REPORT A - UNLIMITED	
2d. DECLASSIFICATION/DOWNGRADING SCHEDULE			
4. PERFORMING ORGANIZATION REPORT NUMBER(S)		5. MONITORING ORGANIZATION REPORT NUMBER(S)	
6a. NAME OF PERFORMING ORGANIZATION SPARTA, INC.	6b. OFFICE SYMBOL (If applicable)	7a. NAME OF MONITORING ORGANIZATION DCA DEFENSE COMMUNICATIONS ENGINEERING CENTER	
6c. ADDRESS (City, State, and ZIP Code) 7926 JONES BRANCH DRIVE SUITE 1070 MCLEAN, VA 22102		7b. ADDRESS (City, State, and ZIP Code) 1860 WIEHLE AVENUE RESTON, VIRGINIA 22090	
8a. NAME OF FUNDING/SPONSORING ORGANIZATION DCEC	8b. OFFICE SYMBOL (If applicable) R640	9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER DCA100-87-C-0095	
8c. ADDRESS (City, State, and ZIP Code) 1860 WIEHLE AVENUE RESTON, VIRGINIA 22090		10. SOURCE OF FUNDING NUMBERS	
		PROGRAM ELEMENT NO.	PROJECT NO.
		TASK NO.	WORK UNIT ACCESSION NO.
11. TITLE (Include Security Classification) DDN Trusted Guard Gateway Phase 2 Report, TRUSTED GUARD GATEWAY (TGG) TECHNOLOGY ASSESSMENT			
12. PERSONAL AUTHOR(S) SPARTA, INC. (D. Solo, C. Eldridge, T. Rivers, G. Ellingswood and J. Mainzer)			
13a. TYPE OF REPORT -BRIEF FINAL	13b. TIME COVERED FROM 6/10/88 TO 2/2/89	14. DATE OF REPORT (Year, Month, Day) -88-2-16	15. PAGE COUNT 101
16. SUPPLEMENTARY NOTATION			
17. COSATI CODES		18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number) COMPUTER NETWORKS, DEFENSE DATA NETWORK. INTERNETWORKING, GATEWAY TRUSTED OPERATING SYSTEMS, VENDORS	
19. ABSTRACT (Continue on reverse if necessary and identify by block number) This document provides to the Defense Communications Agency (DCA) the results of a technology assessment effort made for the purpose of determining the acquisition and/or development potential of a specialized gateway that would interconnect communities with differing security characteristics (such as allowing multilevel secure, classified hosts operating at the unclassified level to communicate with hosts in the unclassified segment.) Section 2.0 of this report reviews the Trusted Guard Gateway (TGG) definition and requirements as outlined in Phase 1 of this effort. Section 3.0 details the Technology Assessment approach and methodology and provides profiles of the various trusted operating system and gateway vendors surveyed. Section 4.0 presents an overview of an alternative approach to the TGG in the form of an application relay.			
20. DISTRIBUTION/AVAILABILITY OF ABSTRACT <input checked="" type="checkbox"/> UNCLASSIFIED/UNLIMITED <input type="checkbox"/> SAME AS RPT <input type="checkbox"/> DTIC USERS		21. ABSTRACT SECURITY CLASSIFICATION UNCLASSIFIED	
22a. NAME OF RESPONSIBLE INDIVIDUAL MR. J. STEVE LLOYD		22b. TELEPHONE (Include Area Code) 703)437-2175	22c. OFFICE SYMBOL R640

TABLE OF CONTENTS

EXECUTIVE SUMMARY.....	1
1.0 INTRODUCTION	3
2.0 TRUSTED GUARD GATEWAY (TGG) DESCRIPTION.....	5
2.1 TGG Definition.....	5
2.1.1 Role of the TGG in the DDN	5
2.1.2 TGG Detailed Requirements	8
2.2 TGG Issues.....	11
2.2.1 Gateway Operation	12
2.2.1.1 Peer Gateway Relationships	12
2.2.1.2 MLS Gateway Routing.....	12
2.2.2 OSI Protocol Support.....	14
2.2.3 Monitoring and Control.....	15
2.3 Shared Responsibility for Protection.....	16
2.3.1 Limitations of TGG Protection.....	16
2.3.2 Host Responsibility for Protection	17
2.4 Backdoor Connections.....	18
2.4.1 Overview	18
2.4.2 Detailed Description	18
2.4.2.1 Types of Backdoor Connection	19
2.4.2.3 Possible Techniques.....	20
2.4.3 Objectives.....	22
2.5 User Requirements Update	22
2.5.1 URDB Consultation	23
2.5.2 Follow-up Contacts.....	24
3.0 TECHNOLOGY ASSESSMENT	26
3.1 Overview and Methodology	26
3.1.1 Assessment Process	26
3.1.2 Assessment Criteria.....	27
3.1.2.1 Protocols.....	28
3.1.2.2 Services	29
3.1.2.3 Certification.....	29
3.1.2.4 Performance.....	29
3.1.2.5 Cost.....	29
3.1.2.6 Schedule.....	30
3.1.2.7 Other	30

3.2	Trusted Operating System Technology	30
3.2.1	Critical Criteria	30
3.2.1.1	Remote Management Protocol Fundamentals.....	31
3.2.1.2	Services Variations.....	31
3.2.1.3	Trusted Operating System Performance Calculation	32
3.2.2	Trusted Operating System Survey	33
3.2.2.1	AT & T System V/MLS	35
3.2.2.2	BiiN/OS	38
3.2.2.3	DEC SEVMS	42
3.2.2.4	Gemini GEMSOS	44
3.2.2.5	Gould UTX/32S.....	47
3.2.2.6	Honeywell SCOMP Trusted Operating Program.....	51
3.2.2.7	Honeywell Secure UNIX.....	54
3.2.2.8	IBM Secure XENIX.....	57
3.2.2.9	SunOS MLS.....	60
3.2.3	Trusted Operating System Assessment Results.....	62
3.2.3.1	General Conclusions.....	62
3.2.3.2	Product Recommendations.....	63
3.2.4	Evolution To A TGG.....	63
3.2.4.1	Development Approaches.....	64
3.2.4.2	Development Level of Effort and Cost.....	67
3.2.4.3	Evolution Summary	70
3.3	Gateway Technology	70
3.3.1	Vendor Interviews	71
3.3.2	Criteria.....	74
3.3.2.1	Schedule.....	74
3.3.2.2	Protocols.....	75
3.3.2.3	Network Management	75
3.3.2.4	Special Services	75
3.3.2.5	Performance.....	76
3.3.2.6	Cost.....	76
3.3.3	Gateway Vendor Survey	76
3.3.3.1	Proteon	77
3.3.3.2	CISCO.....	79
3.3.3.3	Ford Aerospace Corporation (FAC) Multinet Gateway (MNG).....	81

3.3.3.4	Bolt, Beranek & Newman (BBN) Butterfly Mailbridge.....	83
3.3.3.5	Other gateway Vendors	85
3.3.4	Summary.....	85
3.3.4.1	Scarcity of Cross-qualified Vendors	85
3.3.4.2	Need for Partnerships and Roles.....	86
3.3.4.3	TGG Development Estimates.....	86
3.3.4.4	Gateway Operating System Security	88
3.4	Technology Assessment Summary	92
3.4.1	Assessment Overview.....	92
3.4.2	Conclusions.....	93
4.0	APPLICATION RELAY APPROACH	96
4.1	Introduction.....	96
4.2	Application Relay Security Services	97
4.2.1	Authentication.....	97
4.2.2	Application Specific Access Control	98
4.2.3	Labeling.....	98
4.2.4	Guard Functions	99
4.3	Implications of the TAR Architecture	99
4.4	Technology Assessment.....	100
4.5	TAR Summary.....	101



Accession For	
NTIS	CRA&I
DTIC	TAB
Unannounced	
Justification	
By	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A-1	

EXECUTIVE SUMMARY

This report provides to the Defense Communication Agency (DCA) the results of a technology assessment and continued analysis of a Trusted Guard Gateway (TGG). A Trusted Guard Gateway would interconnect communities with different security characteristics (such as allowing multilevel secure hosts on the DISNET operating at the unclassified level to communicate with hosts on the MILNET). This effort involved surveying trusted operating system vendors and commercial gateway vendors and evaluating alternatives for acquiring and deploying a TGG.

The first phase of this program emphasized the user requirements analysis and TGG definition. Certain issues remained relating to gateway security concerns, an evolution to ISO protocols, and an approach for secure remote monitoring and control. In addition, due to recent events associated with the internet worm (or virus) incident, the desire has emerged to provide a finer granularity of access control than that currently defined for the TGG. In order to realize this desire and to provide a more robust firewall, an alternative architecture for the TGG has been developed. Rather than an IP gateway with enhanced access control services, this alternative is a Trusted Application Relay (TAR). This report presents a resolution of the phase 1 open issues and an initial definition of a TAR architecture.

This report covers phase 2 of the TGG program, which is primarily a technology assessment. The primary emphasis of this phase of the Trusted Guard Gateway program is to consider technology bases that might be suitable for the acquisition of TGGs. This report describes our assessment of the two major technology bases: trusted operating systems and commercial gateway products. Based on the initial assessment, providing gateway functionality on top of one of the trusted operating system bases is likely to be less risky than certifying an existing gateway product. The Multinet Gateway (MNG) is the only example of a current trusted gateway product. As the only existing trusted gateway, the MNG serves as a basis against which other candidates may be measured. The complementary advantages of trusted operating system vendors (e.g., cost, certification) and gateway vendors (e.g., functionality, network experience) makes teaming or cooperative options attractive. Interest has been expressed by vendors, particularly in the gateway community, in a joint effort to meet trusted gateway requirements.

In addition to a technology assessment, this report refines the definition of the TGG developed in phase 1. It also presents a possible alternative approach if stronger protection of hosts is required. The technology assessment presents an extensive list of options for the acquisition of a TGG or TAR and measures those options against the TGG requirements. The report also provides rough estimates for the costs associated with developing a TGG based on different technologies. The intent of this report is to present the information with which options may be evaluated. Any eventual

decision on how to proceed with a TGG must reflect a balancing of factors involving cost, performance, functionality, risk, and security.

1.0 INTRODUCTION

The purpose of this report is to provide to the Defense Communication Agency (DCA) the results of a technology assessment and continued analysis effort relating to the consideration of a Trusted Guard Gateway (TGG) that would interconnect communities with different security characteristics (such as allowing multilevel secure hosts on the DISNET operating at the unclassified level to communicate with hosts on the MILNET). This effort consisted of surveying trusted operating system vendors and commercial gateway vendors and evaluating alternatives for acquiring and deploying a TGG.

Section 2 of this report updates the description of the TGG developed in phase 1. Three principal issues concerning the definition of the TGG are resolved. First, the TGG, in addition to providing security services, also has to interact as a peer with other internet gateways. The definition of the TGG must describe the nature of the interaction and must address the security implications of such interaction. Second, DoD has directed a migration to ISO standard protocols. Since the TGG was defined to operate in the TCP/IP protocol environment, the impact of such a transition to the ISO environment needed to be addressed. Third, the TGG must be remotely managed. This management includes the secure monitoring and control of its gateway functionality and the remote maintenance of the security critical information including the access control database.

Section 2 also discusses the shared responsibility for protection between hosts and the network. The motivation for the TGG is to provide a firewall between segments or communities within the DDN. This firewall is aimed primarily at protecting hosts within the more sensitive environment from hosts in the less trusted environment. The network should provide reasonable constraints on the traffic it passes, some guarantees about the information it provides, and should support the actions taken by hosts. It is the responsibility of the hosts to further insure that information that they receive is properly handled within the host, that proper controls on access and privileges are maintained, and that their implementation is sound. In order to provide the overall protection desired for hosts on the DDN with access to other segments, both the network and the host must accept portions of the responsibility. In order to support the policies enforced by the TGG, intersegment paths which are not provided by TGGs, backdoor connections, must be controlled. This section discusses techniques for the detection of backdoor connections.

Section 3 describes our assessment of the two major technology bases required for a TGG: trusted operating systems and commercial gateway products. While the same general process was followed for all the candidate products, the assessment was broadly divided into these two categories. The first category, trusted operating system products, consists of general purpose

operating systems being evaluated under the NCSC commercial product evaluation process. In general, communication networking support is a secondary priority in these types of systems. As a consequence, the assessment concentrated on how these products could be adapted to perform a gateway role. The second category consisted of both trusted and untrusted gateway products. The vendors primarily have products that currently serve as gateways in the internet environment. These products tend to be relatively close to meeting the functional requirements for the TGG; however, certification is a major concern and was emphasized in the assessment.

For both categories, an overall impression of the applicability of the product is presented. This overall impression summarizes how the candidate option maps against the criteria and the perceived magnitude of effort required to enhance the product to meet the full suite of requirements. Factors including cost, schedule, and risk are considered in forming this assessment. Estimates of costs for developing a TGG for each category are presented along with an overall set of ratings.

Section 4 presents a possible alternative approach for providing intersegment security. In light of the recent events associated with the internet worm (or virus) incident, the desire has emerged to provide a finer granularity of access control than that currently proposed for the TGG. In order to realize this desire and to provide a more robust firewall an alternative architecture for the TGG is being developed. Rather than an IP gateway with enhanced access control services, this alternative is a Trusted Application Relay (TAR). Section 4 contains an initial description for how a TAR might operate.

2.0 TRUSTED GUARD GATEWAY (TGG) DESCRIPTION

The first phase of this effort concentrated on the collection and analysis of requirements for a TGG and the definition of the requirements for such a device to operate within the DDN. Along with the basic requirements, a number of issues surfaced requiring further analysis. This section summarizes the requirements developed in phase 1, reviews the open issues concerning the TGG, and presents resolutions for those issues.

In addition, this section discusses the limitations of what a TGG can accomplish within the DDN. The TGG on its own cannot satisfy all the requirements for the protection of hosts. Providing security for host to host operation is a shared responsibility between the hosts and the network. The network should provide some protection and access control as to what connectivity is supported, and should support the labeling that allows the hosts to make decisions on what interoperation it will allow. Beyond these provisions, the host has a responsibility for providing the balance of application data security.

2.1 TGG Definition

In the 1990 time frame and beyond, the DDN will require TGGs to securely allow limited, controlled communications between segments of the DDN operating at different levels of trust or at different security levels. TGGs have been described in DCA plans for the growth and evolution of the DDN.¹²³ The following section describes the planned role of TGGs in the DDN described in "DDN Evolution of Security Services"².

2.1.1 Role of the TGG in the DDN

The DDN evolution is greatly affected by the diverse security and operational requirements of network subscribers. The DDN security architecture evolution has addressed specific needs of classified subscriber communities while maintaining economical approaches to network development and operation. Currently, the DDN consists of multiple physically distinct segments based upon these differing subscriber needs. ARPANET supports continuing access to research and development activities and access by a very wide scientific and academic community. Although the ARPANET is being dissolved, the equivalent collection of networks will remain. MILNET supports unclassified operational needs of DoD agencies. The needs of classified subscribers segregate more naturally into communities

¹ DCA, "DDN Management Engineering Plan"

² DCA, "DDN Evolution of Security Services, 1986 - 1992"

³ DCA, "DDN Subscriber Guide to Security Services, 1986-1992"

of interest. Defense Secure Network 1 (DSNET1) currently serves classified subscribers at the GENSER level. DSNET2 and DSNET3 serve the WWMCCS and SCI communities respectively. While plans are still in the works for a single integrated classified network, current reference to this (i.e., DISNET) translates into DSNETs 1, 2 and 3.

The DDN Security Services Evolution document addresses six types of security services, how they are supported for unclassified subscribers and for classified subscribers, and the security evolution of DDN elements (including policies, procedures and architectural elements). The six security services are:

1. Data confidentiality: mechanisms to prevent unauthorized disclosure of data;
2. Data integrity: mechanisms to prevent unauthorized modification of data;
3. Identification, Authentication and Access Control;
4. Data origin authentication;
5. Non-repudiation: mechanisms to certify to the sender that data were received; and
6. Availability: mechanisms to support assured service of DDN and subscriber resources.

The same document presents the TGG as:

- a support for data confidentiality for both classified (p. 21) and unclassified subscribers (p. 13);
- a support for identification, authentication and access control for both unclassified (p. 17) and classified (p. 25) subscribers; and
- a support for data origin authentication services for classified subscribers (p. 27).

The document describes the TGG as providing interoperability between communities of different security levels (p. 53), different trust levels (p. 39) and between open and closed communities. It does not envision a TGG role in non-repudiation or in data integrity. While a TGG role in availability is not explicitly discussed, Sections 2.3 and 4.4 outline TGG relevance for assuring service.

The DDN Security Services Evolution document states that DCA's plans are to provide for both security and interoperability among DDN segments and between open and closed DDN communities. This report on user requirements and detailed technical descriptions of the TGG supports the execution of these plans.

The evolution of the DDN architecture has taken place in concert with NSA. The INFOSEC organization of NSA has contributed and reviewed

architecture material to support the provision of comprehensive security. The NSA activities have included the development of security systems for individual DDN segments (e.g., BLACKER for DISNET) and the definition, review, and approval of requirements for security across all segments. These requirements have specified mechanisms and levels of trust appropriate for network elements, for hosts, and for the interconnection of segments.¹

Interconnection requirements define functions and assurance levels such that connection to a less secure segment does not compromise the security of the more secure segment. These requirements are addressed in detail for both ARPANET/MILNET and MILNET/DISNET scenarios and directly apply to TGG operation. These requirements dictate the restriction of applications (e.g., mail and file transfer), the prevention of flooding, the labeling of data, and suggest appropriate assurance levels (i.e., B2 for MILNET/DISNET and B1 for ARPANET/MILNET). These requirements serve as a basis for TGG security requirements. Actual requirements reflect a balance between these statements and the operational needs determined in the user survey. These actual requirements are presented in the next section. The resulting compromise must be evaluated for the security provided in the overall DDN.

¹
APR 86

NSA: "INFOSEC Review of the DDN Security Architecture", CSC-TR-26-86, 4




Figure 2-1 TGG Roles Among DDN Segments

2.1.2 TGG Detailed Requirements

The first phase of this program and the associated report¹ discussed issues and presented alternatives associated with defining TGG requirements. The recommended approach balances security requirements, operational requirements, feasibility, and flexibility. The particular functions invoked in a given deployed TGG will depend on static configuration information and dynamic access control tables. This flexibility will allow a common TGG to be used in all scenarios and for the access policy and security functionality of the TGG to evolve with the overall DDN policy and the DDN security posture (i.e., other security systems, certified hosts, certified network components).

¹ SPARTA, "Trusted Guard Gateway (TGG) Requirements Analysis and Detailed Description", 10 May 1988

Based on the user survey, the selection of permitted applications (e.g., mail, file transfer, etc.) must be particularly flexible. While the default requirements may be relatively restrictive, the TGG must be able to accommodate custom applications and interactive traffic for some subset of users.

The following paragraphs define the high level baseline requirements for the TGG:

1. The TGG must be capable of performing all the standard functions expected of an IP gateway for the DDN.
2. The TGG must support network interfaces for ARPANET, MILNET, and DISNET as well as for the BLACKER Front End.
3. The TGG must label all datagrams passing from a low side segment to a high side segment. The label will be placed in an extended IPSO field (IP Option 133) as described below. Because network segments are concatenated, a TGG may receive a datagram already labeled by another TGG. Depending on configuration information, a datagram already labeled will be:
 - a) audited and discarded, with a message sent to the source,
 - b) passed without modification subject to other access control checks, or
 - c) have the TGG count incremented.
4. The TGG must be able to limit the flow of datagrams from a low side segment into a high side segment. This limitation will be based upon a TGG's configuration. The TGG will keep counters for any limitations and discard datagrams exceeding those limits. Thresholds may be set for:
 - a) total number of datagrams across an interface,
 - b) total number of datagrams from a particular source address,
 - c) total number of TCP connection requests (as reflected by SYN TPDUs), or
 - d) number of access control rejections (this will shut down an interface).
5. The TGG must enforce access control rules on every datagram. The TGG will maintain an access control database that defines permitted and restricted datagrams by interface and by source address. The following fields are potentially subjected to access control checks:

- a) IP source network number, restricting where traffic can originate;
- b) IP source and destination addresses, restricting which host pairs can communicate;
- c) IP protocol field, limiting the transport protocols that are allowed (such as prohibiting UDP applications);
- d) TCP port field, identifying the application being used and restricting the use of mail, file transfer, virtual terminal, etc. This field also indicates the originator of the application and can restrict application direction. TCP port access control rules are maintained on an interface, and on a source and destination address basis; and
- e) IP security options as discussed in the requirement for labeling.

6. The TGG must support secure remote monitoring and control. The formats, mechanisms, and protocols for monitoring and control should be consistent with evolving standards for network management, especially those being developed for DDN. Monitoring and control functions include:

- a) status and health reporting,
- b) remote tests and diagnostics,
- c) security audit reports,
- d) control of TGG configuration parameters identified above, and
- e) control of the access control database described above.


7. Assurance must be provided that the TGG securely performs these functions, protects data, and cannot be circumvented. This assurance is provided in part by COMPUSEC Certification. Based on the environment, the DDN plans, performance requirements, and technical availability, the TGG should be certified to a B1 level. The selection of the B1 level and its interpretation are discussed in the phase 1 report.

8. In the baseline TGG configuration the following options are selected:

- a) label datagrams with extended IPSO and pass datagrams that are already labeled,
- b) no flow limitations across any interfaces, and

- c) access rules are enforced on IP protocol field and TCP port field for all datagrams across an interface. Only TCP is allowed and only mail and file transfer applications are supported.

The requirements above include the need to provide labeling of datagrams originating on the low side using the extended IPSO. The format for the use of that option is shown in the figure below.



The ID for the use of the extended IPSO by TGG's is to be assigned by DCA. The length field is 4 when used in DoD IP (when used in OSI CLNP, the length field is not currently specified). The data field is one byte and includes three bits of count and five bits reserved. The reserved bits are set to zero. The count bits indicate the number of TGG's which have processed the datagram. The use of the count field depends on configuration options as described above and is set to one by the first TGG to process the datagram.

These requirements are based upon a TGG operating as an internet relay, (i.e., an IP gateway with value added security services). As a consequence of the recent events associated with the internet worm attack, the issue of defining adequate security services for the TGG has resurfaced. In this light, we have also considered the requirements for a TGG operating as an application relay. A discussion of the limitations of a TGG is presented in section 2.3 while a discussion of a Trusted Application Relay (TAR) is presented in section 4.

2.2 TGG Issues

In the first phase of this program, the requirements described above were developed. In addition to those requirements, certain issues pertinent to the operation of a TGG remained unresolved. This section summarizes the open issues and describes their resolution.

2.2.1 Gateway Operation

In addition to performing its specialized functions as a firewall in the DDN, the TGG must act as a gateway within the DDN system. This requirement means that the TGG must interact with untrusted gateways as well as trusted and untrusted hosts. The open issue identified in phase 1 involves the security implications of these interactions and the protocols to be used. Resolution of those issues is presented below.

2.2.1.1 Peer Gateway Relationships

The issue of peer gateway interaction is driven by the evolving standards within the DDN and the overall internet. In order to effectively operate, the TGG must comply with existing standards. Given this constraint, the major remaining question is the TGG's role within the autonomous system structure. Based on current plans within the DDN, there is no need for the TGG to act as one of the core routing gateways. Consequently, the TGG will exchange reachability information with peer gateways with the appropriate exterior gateway protocol. For the TGG timeframe, this is expected to be EGP¹. Similarly, the interaction with hosts will be through ICMP.

These decisions may need to be revisited as the standards evolve. Of particular concern is a transition to OSI standards and the evolving Intermediate System to Intermediate System (IS-IS) and End System to Intermediate System (ES-IS) standards. In any event, these gateway protocols may need to be supplemented, especially by authentication, to accommodate the security issues described below.

2.2.1.2 MLS Gateway Routing

An extensive amount of work has been performed recently in advancing the design and performance of routing algorithms. Much of this effort has assumed an essentially flat topology. A flat topology is an environment where transitivity applies to connectivity: if gateway A can communicate with gateway B, and gateway B with gateway C, then A and C can communicate. In an environment where gateways and the networks they connect have associated ranges of classified data, this assumption need not hold. Disjoint, partially overlapping, and equivalent security ranges may exist for a wide range of gateways.

¹ IETF, "IDEA0009: Exterior Gateway Protocol, Version 3," Feb 1988

MLS routing is further complicated by the fact that topology information may itself be sensitive. This sensitivity derives both from an intrinsic importance attached to information concerning the existence and connectivity of the networks and from the possibility that a gateway may intentionally or inadvertently include user traffic in routing updates. The association of sensitivity labels with routing information and updates complicates the generation of routing updates sent to other gateways and of control information sent to hosts such as redirects or other ICMP messages. In such an environment, any message sent by a gateway must be a function of the sensitivity of the information on which the message is based, on the clearance of intermediate entities, and on the clearance of the entity to which the gateway is communicating. This means that a gateway may advertise substantially different views of the internet to otherwise equivalent gateways on an attached network.

In addition to affecting information propagated through gateway control messages, the multilevel nature of the networks influences the routes that are calculated for individual messages. This is a specific instance of the notions of policy based routing that are being described by Dr. David Clark¹ and others. The problem described in the policy routing work is further complicated by (or possibly simplified by) the situation described above which limits the propagation of routing information.

Another issue to take into account in MLS gateway routing is the possibility of denying service or degrading performance through the distribution of malicious or intentionally incorrect routing information. Some attention has also been given to the design of routing algorithms and protocols that minimize the damage bad information can cause. EGP and the autonomous system concept in part address this issue. However, these approaches do not address the concerns of hosts masquerading as legitimate gateways. The concerns about bad routing information need to be addressed both through the inclusion of authentication information in gateway control messages and in the design of routing algorithms that minimize the damage that may be caused by a malicious, legitimate gateway.

The problems of authentication are increased if the gateway community is a large, diverse, heterogeneous collection under a variety of administrative controls. Many authentication schemes require a single central authority to act as the manager for system credentials. Other distributed schemes offer more flexibility at the cost of reduced confidence in the authentication provided. Future gateway routing algorithm standards will need to consider a variety of schemes for use in gateway authentication.

These aspects of an MLS internet must be taken into account in designing a routing protocol and gateway processing description that can

¹ "Policy Routing in Internet Protocols," Dr. David Clark, DRAFT May 1988

provide the required services. The design must address what information needs to be passed between gateways and between gateways and hosts to support an MLS environment, how the information passed is authenticated, how that information must be acted on within gateways, how routes should be calculated, and how the sensitivity of the information involved is handled. The solution must accommodate the management of these gateways as well.

For the TGG, the scenarios identified in the DDN allow for some simplifying assumptions. These are based on the fact that the TGG is handling single level data. By relying on the multilevel secure gateways in the DISNET to provide whatever filtering is necessary, the TGG does not have to implement special filtering mechanisms. The need for gateway authentication remains important in any case. This problem cannot be solved unilaterally within the TGG program, but rather must be addressed as part of the evolution of the gateway to gateway and gateway to host protocols (i.e., IS-IS and IS-ES Protocols).

2.2.2 OSI Protocol Support

Based on directives from OSD¹ and on the recently published DDN OSI Implementation Strategy², the DDN expects to move towards OSI protocols in the 1990's, initially as costandards and eventually as the primary standards. The open issue addressed here is the consideration of the impact of a transition to OSI upon TGG functionality. The likelihood of vendors' support for OSI protocols was examined during the technology assessment.

The first area where the OSI transition will impact the TGG is in the evolution to the OSI Connectionless Network Protocol (CLNP or ISO IP). From a functional standpoint, this poses no obstacles to TGG operation. The addressing and other header fields are consistent with the TGG definition and the IP security option defined for DoD IP was defined to be a valid OSI option as well. In discussions with vendors, there is near universal agreement that the migration will occur as the market and policy dictates.

The second area involves gateway and management protocols, also referenced in sections 2.2.1 and 2.2.3. In general, the OSI versions of these protocols are among a set of evolving and competing standards. In defining the TGG requirements it is assumed that the TGG will have to comply with whatever management and gateway protocols emerge from the current process. The OSI suite is as acceptable a choice as others. Further, the OSI standards are likely to include placeholders for authentication.

¹ ASD/C3I Memo, "Open Systems Interconnection Protocols," July 1987

² MITRE, "The Department of Defense Open Systems Interconnection (OSI) Implementation Strategy," DRAFT May 1988

The third area is demultiplexing. Within the OSI framework, the information available at a given protocol level indicates a service access point at the next level. Given a single protocol at each level, the only information provided is the next layer protocol. The strict layering dictates then that at the network layer, the only information is that TP is the next layer transport protocol; then at the transport layer, that the next layer is session; and so forth. In addition, there is no guarantee as to how PDU boundaries will align. This is contrasted with the TCP/IP case where examination of the TCP port field indicates both the application and the initiator. In order to achieve like functionality in the OSI environment, significantly more header processing is necessary. Since current approaches for demultiplexing do not guarantee the encoding of application type in transport and session level identifiers, an OSI TGG would need to either have significantly higher processing power or would provide significantly lower performance.

2.2.3 Monitoring and Control

The open issue associated with monitoring and control for the TGG includes two aspects. The first aspect is that as a gateway in the internet system, the TGG must be managed along with the other major gateways. This requires the reporting of health and status information and limited maintenance functions. The second aspect is the management of the security role of the TGG. This is primarily concerned with the remote management of the access control information.

A goal for the TGG is to use standard monitoring and control protocols for both of these functions. The alternative of a specialized TGG control center is a possible short term solution but is viewed as undesirable. Evolving standards within the internet community seem capable of supporting the needed functions. This scenario relies on the authentication of management traffic and on an appropriate level of trust in the command center. In assessing technology options for the TGG, the lack of any secure remote management capability became very clear. Since no short term, off the shelf solution is available, the choice of authenticated standard protocols is appropriate.

An area of concern relates to the sensitivity of management information. This is based both on the intrinsic sensitivity of the information reported to a monitoring center and on the possibility that user traffic might be reflected in management traffic. The concern about user information being reflected in management traffic relates to cases where both the protocol and management function are untrusted. In such cases, the data sent to a management function, and in turn relayed to a monitoring center, may be user data. If either function is trusted, then the data passed may be checked to insure that no user data, or only very limited user data, is released through such a covert channel with the management traffic consequently treated at a network wide "management level."

The previous paragraph suggests that management information should be treated as the highest level of the information processed by the device being managed. For the TGG in the ARPANET/MILNET and MILNET/DISNET scenarios, this is UNCLASSIFIED. The TGG analysis assumes that the monitoring center will be trusted not to disclose any information to the TGG that is more sensitive than the information that the TGG normally handles. This may be achieved through either a trusted monitoring center or separate community of interest monitoring centers.

2.3 *Shared Responsibility for Protection*

The motivation for the TGG is to provide a firewall between segments or communities within the DDN. This firewall is aimed primarily at protecting hosts within the more sensitive environment from hosts in the less trusted environment. Providing the protection for these hosts is a shared responsibility between the hosts themselves and the network. The network should provide reasonable constraints on the traffic it passes, some guarantees about the information it provides, and should support the actions taken by hosts. It is the responsibility of the hosts to further insure that information they receive is properly handled within the host, that proper controls on access and privileges are maintained, and that their implementation is sound.

In order to provide the overall protection desired for hosts on the DDN with access to other segments, both the network and the host must accept portions of the responsibility. The requirements defined above are an appropriate set of functions for the network to provide consistent with general network services. Section 4, the description of a Trusted Application Relay (TAR), presents an alternative where the network undertakes a much larger share of the responsibility. This section describes some of the limitations on what the currently defined TGG can achieve and some of the guidelines for hosts utilizing the services of the TGG.

2.3.1 Limitations of TGG Protection

The primary security functions of the TGG are to provide access control on a host pair and application basis and to label datagrams originating in the low side as being of "suspicious origin". These functions provide a host with a label that allows it to restrict processing of datagrams and with a guarantee that low side messages will only arrive on certain well known ports.

A TGG that examines only IP and TCP header information does not provide any guarantees that the actual application corresponds to the application assigned to the well known port, nor does it provide any help in closing weaknesses or vulnerabilities within applications. A pair of conspiring hosts could readily exchange virtual terminal traffic between the SMTP port numbers. The only possible way the TGG might eventually detect this without parsing application header information is by observing the

substantially different traffic patterns of virtual terminal service as contrasted with mail service. Further, if weaknesses in the SMTP specifications or implementations exist, the TGG provides no help.

Despite these limitations, the TGG provides significantly improved protection for communication services provided between segments by the DDN. This alternative is likely to be attractive to subscribers particularly if the alternative is to sever all connections. The experience in the past has shown that if no direct service is provided (and possibly even if it is), backdoor connections will be established. These connections represent hosts, LANs, or collections of networks that are homed to multiple segments (e.g., to both the ARPANET and MILNET). These connections are usually completely uncontrolled. If no direct connection is provided, then backdoor connections are more likely to come into existence in order to support the user's need to satisfy his mission requirements. Providing the TGG class of service with the flexibility to meet mission requirements is a substantial improvement over the uncontrolled connections that are a likely consequence of providing no service at all.

The nature of these backdoor connections points out the importance of managing such connections in order to provide an effective TGG service. Where such connections are authorized, a gateway or host on the path must implement the same features as provided by a TGG. All other connections need to be actively discouraged and terminated where possible. A discussion of approaches for detecting backdoor connections is provided in section 2.4.

2.3.2 Host Responsibility for Protection

The complement of the issues presented above is that hosts that are authorized to use the TGG service will need to comply with guidelines on their operation. These guidelines can be voluntary based on warnings about the alternatives or can be enforced through conformance testing. The examples here are representative of some of the issues that might be addressed in a set of host guidelines, but are not intended to be comprehensive.

Hosts should be capable of intelligently processing the label that indicates a datagram is from a less trusted segment. That information could trigger host specific checking or containment and may be passed on to the human user. The implementations of standard protocols should conform to guidelines established for those applications covering features to be supported and methods for supporting them. The host should implement strong password management schemes designed to resist typical password attacks. The host should implement appropriate discretionary access controls so as to limit the exposure of information and of system resources.

When combined with reasonable guidelines for hosts authorized to communicate to less trusted segments, the TGG can provide an effective level of overall protection across the DDN system. By offering access to TGG

services only to those hosts that comply with the guidelines, by limiting the applications used, and by labeling datagrams, the TGG provides a significant assist to the user information protection process.

2.4 Backdoor Connections

During our analysis of trusted gateways, the importance of detecting and controlling backdoor connections became apparent. Trusted gateways (or other trusted relays) act as firewalls between segments of the Internet with different security characteristics. The firewalls provide enhanced access control and limit intersegment operation. Controlling intersegment communication requires that a common set of rules is enforced for all traffic and on every path between segments. The trusted gateways provide direct connections between segments and are intended to carry all intersegment traffic. Other paths provided by hosts, networks, or collections of network connected to both segments may exist. The existence of these "backdoor" connections violates the assumptions needed to provide the most effective trusted gateway service.

2.4.1 Overview

Backdoor connections tend to be uncontrolled paths between segments that allow arbitrary interoperability. In order to support the trusted gateway policies, and to provide the most effective trusted gateway protection, these backdoor connections must either be controlled or eliminated. Though the process of reducing these backdoor connections may be a gradual one, the direct intersegment connections will still carry the majority of intersegment traffic. Consequently, the control and management of that portion of the intersegment traffic is a positive step in providing overall protection of the hosts in the DDN. Eliminating backdoors is viewed as a means for enhancing the utility of trusted gateways, not as an absolutely necessary prerequisite for deploying them.

One of the most important reasons for maintaining connections between segments is to provide automated support for a reaction when problems do occur. The recent worm incident has highlighted the need for coordination and for the distribution of information in order to best respond to and deter such incidents. Overall protection can be enhanced by continuing to provide some intersegment service even when the security on that service is less than perfect.

In addition to the use of administrative sanctions against backdoor connections, the operator of the MILNET (and other similar situations) needs an ability to detect and shut down such connections. This ability represents part of the enforcement mechanism for an administrative prohibition on intersegment connections. Approaches for the detection and reduction of backdoor connections involve a combination of technical and administrative techniques.

2.4.2 Detailed Description

An approach to detecting backdoor connections is based on the fact that the detection of these paths is closely related to the process the Internet (i.e. hosts, gateways, etc.) must follow to route information to its destination. By exploiting the information used by routing agents (e.g., internet gateways, application relays, hosts), we can find the backdoor connections we are concerned about if they exist. This section will describe the types of backdoor connection we are addressing and the types of techniques which may be used in detecting those connections.

2.4.2.1 Types of Backdoor Connection




Figure 2-2: Backdoor Connection Categories

Backdoor connections can be broadly divided into four categories to be addressed:

- 1) Direct gateway connections;
- 2) Indirect gateway connections;
- 3) Dual homed hosts;
 - A) No general relay services
 - B) General application relay services provided
 - C) Deliberately deceptive hosts
- 4) Dial up lines.

Category 3, dual homed hosts, is subdivided based on the mode of operation of the host. Dual homed hosts can provide differing types of relay services. In some cases, covered under 3B, these hosts provide generally available, widely publicized relay services (Mail relay hosts between environments are examples of such hosts). This subcategory is the type of dual homed host which is the focus of backdoor detection. A second type, covered under 3A has two separate protocol stacks for each network connection and provides no automated relay capability. These hosts provide service to both segments only for logged in users and are consequently of lesser concern. The third type of dual homed host, covered under 3C, is characterized as actively seeking to avoid detection. While it may provide either class of relay capability described above, it is assumed that such a host is aware of backdoor detection activities and is behaving so as not to trigger any of the techniques discussed below.

Categories 4, 3A, and 3C represent connectivity that, in general, would be difficult to detect, but which is available only to a few limited "subscribers". Without dissemination through Internet reachability protocols, these types of backdoor paths will not be generally known or available. From the standpoint of risk to the MILNET, categories 4 and 3A are equivalent to normal MILNET hosts which may misbehave and, as a result, are covered by standard host rules. Categories 1,2 and 3B are the principal targets of a detection process. These categories involve Internet connectivity between segments either directly or through multiply homed networks and application connectivity through relays providing a regular service.

2.4.2.3 Possible Techniques

This section describes an assortment of possible techniques for detecting backdoors. A large variety of information exists within the Internet which can provide an indication of possible backdoor connections. The operation of the Internet relies on the ability of hosts to discover routes to their desired destination. The duplex nature of network protocols requires that in addition to the source host finding a route to the destination, the destination must be able to find a route back to the source. Some of the techniques discussed here attempt to exploit some of the same sources of information.

Once a catalog of these techniques is collected, they can be further assessed in light of operational characteristics in order to determine their actual utility. The consideration of possible techniques for detecting backdoor connections serves to identify what can be achieved in the control of such connections. In addition, the process of considering alternative techniques and of establishing limits helps to define the administrative procedures needed to make detection more practical and which are needed to supplement technical control approaches. An example of an administrative prohibition is a ban on MILNET hosts doubling as IP gateways.

Techniques may be passive or active. Passive techniques involve the collection and analysis of data available by observing and monitoring the network operation. Active techniques involve the probing of possible paths, the stimulation of the network, the exploration of network topology, and the examination of suspicious information.

Examples of passive information include the examination of the routing tables in both core and EGP gateways, the examination of routing updates exchanged between gateways, the examination of mail headers and other protocol headers, the review of host tables, and possibly the review of bulletin boards and mailing lists.

Routing information provides an indication of connectivity in the Internet. In addition to searching for paths other than those through the official, trusted gateways, examination of routing anomalies can help. Datagrams or routing messages which produce routing errors can be suggestive of the attempted use of backdoor connections. Information in mail headers relating to the forward and return routes can provide indication. An indirect source of information is the explicit and implicit directory services for locating paths. If information in host tables can be examined, that information should indicate a host's perceived route to a destination. Passive information may be obtained through EGP requests, current management protocol requests, or other monitoring techniques.

Examples of active probes include the use of IP record route and loose source route options on ICMP echo messages; sending mail requiring forwarding; and queries of gateways, hosts, and domain/directory servers. The "traceroute" tool developed by Van Jacobson, which is another example of an active probe, explores black holes by progressively incrementing the time to live of datagrams. Dramatic tests could also include the temporary deactivation of the official intersegment relays (e.g., trusted gateways, Mailbridges) or other topology perturbations.

The combination of loose source routing and record route options can be employed to inject datagrams into networks connected to the MILNET to determine how they behave with ARPANET destined datagrams. Probes of gateways looking for redirects are another possible test. Submitting mail to suspected hosts requesting relaying to the ARPANET is a possible means of detecting all types of category three backdoor connection.

One of the issues in any of these techniques is where and how the instrumentation is provided. While certain information is already available through the existing network monitoring and control activities, other information must be obtained through additional collection at nodes already monitored or additional collection nodes. The assessment of these techniques must include the identification of where the information is collected.

Current DDN operational mechanisms may support the management of backdoor connections. The community of interest separation tools at the network and Internet level can provide isolation which both augments the elimination of backdoors and which assists in their detection.

2.4.3 Objectives

The goal of these techniques is to support controlled intersegment communication in the DDN. This communication is primarily the interoperation between Internets based around the MILNET and ARPANET. The TGG provides trusted connectivity between segments. The effectiveness of those approaches is based on the success in controlling other paths between the segments. While complete success in eliminating backdoor connections may not be possible, progress in controlling those paths provides significant improvements in the overall protection afforded hosts on the MILNET.

The expected scenario for controlled intersegment operation involves a number of measures. The first is to deploy appropriate relays between the MILNET and the ARPANET (e.g., Mailbridges, trusted gateways, trusted application relays). The second is to systematically detect and eliminate uncontrolled paths between segments. The third is to improve the techniques for the control and detection of intersegment paths. The fourth is to provide standards and procedures for the approval of multihomed subscribers. This last aspect provides an option for backdoor connections once detected. They may either be eliminated or they may be made to conform to controlled intersegment requirements. This scenario provides a means of meeting user requirements while steadily improving the overall DDN security.

2.5 User Requirements Update

The emphasis during the second phase of this effort has been on the technology assessment, however for continuity's sake, an effort has been made to maintain contact with users identified and surveyed during phase 1. Consequently, some observations and conclusions can be made in the area of users and user requirements.

Phase 1 involved the surveying of major DDN subscriber communities with regard to the need for and utility of TGGs. The aims of the survey were to validate subscriber TGG requirements and to identify the numbers and types of TGGs required. During the process, the User Requirements Data Base

(URDB) was a major factor in finding appropriate contacts responsible for the systems. It also provided a significant amount of statistical data useful in characterizing the DDN and the kinds of networks planning to attach to it. Along with the user assessment, SPARTA provided some conclusions on the usefulness of the User Requirements Data Base as a planning and tracking tool for the evolving DDN. This is another issue revisited briefly for this report.

2.5.1 URDB Consultation

As explained in the final report for phase 1, the URDB is the information repository for subscribers' validated requirements for DDN network connectivity. It contains descriptions of subscriber hosts and terminals; intended sources and destinations for data communications; and data is organized into a collection of relational databases. The survey began with a study of the definitions of the relational database files.

For phase 2, we relied heavily upon the previous work. Extensive use had been made of the URDB in April just prior to the end of the first phase. Access to the URDB lapsed and, due to the administrative requirements of initiating new paperwork and getting it approved by the contracting officer, access was not renewed until late in November of 1988. This turned out to be a tactical error on our part as the URDB has undergone significant changes over the past year. Many of the files accessed by our previously written programs had changed and no longer responded to query. As the technology assessment was the focus of this phase, it was decided to gain whatever new data was possible through the execution of remaining working programs, rather than rewriting programs.

While this situation proved to be disappointing as it did not allow for as in-depth a review as would have been preferred, it does result in a positive note. The same changes that have made the URDB temporarily less useful for our purposes, seem likely to be the kind of changes needed for the URDB to be more useful to users and administrators alike. A review of the files on the system demonstrates that not only are there a larger number of files, but that the files have been refined considerably. This will no doubt result in a more user-friendly situation and may mean that manipulation of files will yield a finer granularity of detail, providing administrators with a more realistic view of what is happening on the DDN.

Some of the runs that were successfully performed confirmed preliminary conclusions made during the first phase of the effort. The DDN is growing rapidly and overall requirements for the communities planning to attach to it are frequently changing and evolving. As to the basic topology of the DDN, no significant change was noted other than that of communities growing larger. Neither have the basic character of system requirements changed significantly, though there is a feeling that these requirements are

becoming better understood by those responsible for systems. This is delineated in more detail in the section below on follow-up contacts.

2.5.2 Follow-up Contacts

In phase 1, the surveyed systems were classified into four categories. (For complete descriptions of the categories as well as an in-depth explanation of the formulas used to turn the survey into actual TGG requirements, the reader is referred to the final report for Phase 1, Trusted Guard Gateway (TGG) Requirements Analysis and Detailed Description; May 10, 1988.) For the purpose of this report it is sufficient to say that the categories were based upon factors including size and character of system, level of definition of needs and requirements, knowledge of systems' administrators, and system mission. Given the rapid growth in the number of systems desiring to connect to the DDN, it would be interesting to pursue these categories to determine any changes. However, for this phase of the contract those systems (i.e., category A systems) that formed the basis for the original determination of requirements for the TGG were the focus.

Category A systems are larger, distributed systems with relatively well defined needs and requirements (e.g., databases requiring access to and updates from field sites). Generally, it was found that the people contacted for these systems had a better understanding of security requirements and criteria for certification than administrators of systems in the other categories. This remained true with the category A systems that responded to follow-up discussions during phase 2.

In general, requirements for these category A systems had changed only in that the requirements themselves had become better defined. Of particular interest is that people in general are becoming more informed about their systems and the requirements of its users. In addition, there seems to be more of a concern with continuity. During the first phase of the effort, one of the things that made it very difficult to gather information about systems and users was that the military system finds people changing positions frequently and, more often than not, new people are not privy to all of what has gone on before. Three systems that we contacted again had undergone such a change, but the predecessors of those now responsible saw fit to pass on more than the usual amount of information about the systems and plans for the DDN. This is definitely an encouraging sign in that there is a growing realization that asking people to start from scratch does not benefit anyone, particular the users of the system.

The brief sojourn back into the user requirements left us with several positive indications. The URDB appears to be becoming an even better tool for DDN administrators to perform tracking and planning and should continue to receive the attention that has brought about the recent improvements. The requirements for the TGG as seen in the user survey in the first phase remain, but are becoming better focused and understood. And

finally, those responsible for large systems on or intending to attach to the DDN are becoming better informed about their systems and user populations.

3.0 TECHNOLOGY ASSESSMENT

3.1 Overview and Methodology

The primary emphasis of this phase of the TGG program is to consider potentially suitable technology bases for the acquisition of TGGs. These bases are commercial or special products that could be augmented to meet the requirements described in section 2. A number of possible paths for acquiring and deploying a TGG exist. First, a secure gateway product could be found requiring few, if any, extensions. Second, an untrusted gateway product could be enhanced and subsequently certified. Third, a trusted operating system product could be selected and enhanced by adding gateway functionality. Fourth, a custom development could be undertaken. The first three are all considered valid options and are explored in this section. The fourth option has not been considered as it represents a less efficient choice due to the duplication of capabilities with existing products.

This section describes our assessment of the two major technology bases needed for a TGG -- trusted operating systems and commercial gateway products. Presented are the methods of gathering information, the assessment criteria, the findings for both technology bases, and overall observations and conclusions.

3.1.1 Assessment Process

The process for conducting the assessment involved four steps. The first step was the development of a list of candidates. The list was based on our familiarity with the community, references in published literature, and recommendations from individuals contacted during the program. The second step was to obtain preliminary information through telephone or mail contacts. Following this initial contact, a more detailed discussion would take place at a meeting between the vendor and SPARTA. These meetings served to both further explain our role and objectives and to assess the applicability of the vendor's product(s). The fourth step was to perform follow up contact on unresolved issues or further questions revealed during the analysis process. These follow ups could be conducted either over the phone or with subsequent face to face meetings.

While the same general process was followed for all the candidate products, the list is broadly divided into two categories. The first category is the trusted operating system products. These are general purpose operating systems being evaluated under the NCSC commercial product evaluation process. Generally, they are systems where communication networking support is a secondary priority. As a consequence, the assessment concentrates on how these products could be adapted to perform a gateway role.

The second category consists of commercial gateway products, both trusted and untrusted. These are vendors who have products that currently serve as gateways in the internet environment. These products tend to be relatively close to meeting the functional requirements for the TGG; however, certification is a major concern and is the primary emphasis of the assessment.

For both categories, an overall impression of the applicability of the product is presented. This overall impression summarizes how the candidate option maps against the criteria and what we believe to be the magnitude of effort required to enhance the product to meet the full suite of requirements. Factors including cost, schedule, and risk are considered in forming this assessment. Where our impression is based on subjective assumptions or on alternative scenarios, those issues are identified.

3.1.2 Assessment Criteria

The requirements for the TGG described in section 2 serve as the basis for the evaluation criteria for technology options. While the list of criteria is the same for both operating systems and gateways, the emphasis between the two vary. For this reason, the following two sections, 3.2 and 3.3, address the criteria in more detail for each category of implementation option. Figure 3-1 lists the criteria applied to candidate TGG options.

TGG IMPLEMENTATION ALTERNATIVE CRITERIA

- PROTOCOLS
 - COMMUNICATION
 - GATEWAY OPERATION
 - MANAGEMENT
- SERVICES
 - ACCESS CONTROL
 - LABELING
 - FLOW LIMITATION
- CERTIFICATION/ASSURANCE
- PERFORMANCE
- COST
 - DEVELOPMENT
 - PRODUCTION
- SCHEDULE
- OTHER

Figure 3-1: TGG Assessment Criteria

3.1.2.1 Protocols

The criteria listed under protocols represent those requirements necessary for the TGG to operate as a gateway in the DDN. The evaluation against these criteria involve the support for protocols that conform with existing published or *de facto* standards for gateway to gateway operation, network interfaces, internet protocols, and management functions. While the host and gateway versions of IP are different, support for a host IP implementation is considered a positive factor since it indicates a favorable environment for IP services.

A gateway IP implementation must support the forwarding of datagrams (relaying datagrams from one lower level interface to another) and must perform fragmentation. A gateways ICMP operation is also different. These gateway protocol requirements are discussed in more detail in the gateway requirements RFC¹. The IP implementations discussed in the section on

¹ "RFC 1009: Requirements for Internet Gateways," Jun 1987

trusted operating systems relate to host IP implementations. Gateway IP implementations are also available for UNIX based operating systems.

3.1.2.2 Services

The criteria listed under services are the value added services performed by a TGG. This is the heart of the TGG functionality intended to provide a firewall between segments of the DDN. The identified services include access control at a host pair and application level, labeling of IP datagrams with a "suspicious origin" label, and the limitation of the flow of datagrams based on configurable thresholds. While the details of all of these services are specialized, the assessment looked for vendor functionality that represented a close approximation of the desired service and that would require minimal changes in producing a TGG. Audit capabilities covering these services are also needed.

3.1.2.3 Certification

The requirements for the TGG indicated the need for a trusted computing base (trusted network base). This criteria category assesses the status of the certification of the candidate product and, for those products not in the certification process, an estimate of the certifiability of the product.

3.1.2.4 Performance

Measurements of the performance of gateways may be described in a number of ways. As a consequence, even for commercial gateway products, advertised values for performance may not be directly comparable. While we have presented numbers for datagrams per second throughput for the candidates, these numbers should be viewed as rough order of magnitude estimates rather than absolute values. For this assessment we are concentrating on a requirement derived for a DDN environment characterized by links at T1 or slower rates. The assessment would produce different results for environments where the TGG was expected to operate at rates of 50-1000 Mbit/sec.

3.1.2.5 Cost

In determining cost, two scenarios are possible. The first is to assume that special development and certification costs are funded or recovered separately and that the cost represents the unit production cost. In this case, the cost is close to that of the current commercial product. In the second scenario, special costs are amortized over the quantity of gateways sold. Because of the small quantity of TGGs (less than 100) and due to the uncertainty in estimating special costs, the figures presented for cost are based on current commercial products. Development cost estimates are presented at the end of sections 3.2 and 3.3 along with unit costs assuming an amortization of those costs.

3.1.2.6 Schedule

The schedule criterion reflects an assessment of how current product schedules meet the TGG schedule. This primarily is based on certification schedules but also reflects the magnitude of the functional enhancements required and the vendor's ability to aggressively pursue those changes. All references to dates in terms of quarters are with respect to calendar years.

3.1.2.7 Other

This category is reserved for additional information that is considered relevant in forming the overall impression.

3.2 Trusted Operating System Technology

This section presents the findings of the trusted operating systems survey and the definition of an approach for the evolution of a trusted operating system product to a TGG. Prior to the actual survey it was expected that certain criteria would prove critical to the utilization of a trusted operating system product in the role of a TGG. These criteria, as identified and analyzed in the following section, form the basis for the survey, assessment, and subsequent conclusions. Based on the overall assessment conclusions, the additional functionality necessary to evolve a trusted operating system product to a TGG is identified. This includes the identification and definition of possible design approaches with associated level of effort and cost information.

3.2.1 Critical Criteria

A subset of the TGG implementation criteria described in section 3.1.2 proves critical to the utilization of a trusted operating system as the foundation for a TGG. The specific criteria are protocols, gateway services, performance, and cost. These criteria are identified in the following paragraphs and detailed in the following sections as appropriate. Further details are provided concerning a remote management protocol, gateway services, and performance.

Due to the general purpose nature of a trusted operating system, communication, gateway operation, and remote management protocols are not extensively supported. This is especially true of gateway operation protocols, which are not supported by any of the candidate trusted operating systems. In support of the remote management functionality, the basic security functionality of a trusted operating system must be expanded to support a network environment. Network security functionality is only now in its early stages of research and development.

Services in the form of access control, flow limitation, audit, and labeling are supported by trusted operating systems, however the interpretation of these services in an operating system is substantially different than in a

gateway. Of concern is the adaptability of these services to a specific application, such as a TGG.

Basic performance is generally acceptable, however the constraints placed on basic performance by the specific TGG application must be addressed. Also, providing a means of evaluating the performance of trusted operating systems in terms of datagrams/second is of concern.

Finally, cost is a factor in a number of trusted operating systems. Although, when compared to a equivalently certified gateway product cost typically becomes of less concern.

3.2.1.1 Remote Management Protocol Fundamentals

The fundamental requirement necessary to support a remote management protocol capability is authentication. This capability would provide a network version of a trusted path that would enable the trusted maintenance of host pair access rights, definition of audit parameters, collection of audit data, and other relevant security actions. Ideally, the secure networking capability would function independently of the various host's hardware/software complement. None of the trusted operating systems surveyed and assessed meet this requirement. This is primarily due to a lack of standards in this area which leaves vendors to either await a standard before developing this functionality or to develop this functionality independently.

Many of the trusted operating systems support a non-secure remote management capability, while a smaller number provide a secure remote management capability within a network made up entirely of their own products. In the following assessments, the ability of each vendor's product to support a remote management capability is identified and detailed.

3.2.1.2 Services Variations

The nature of the required services in the form of access control, flow control, auditing, and labeling for the TGG differs from that implicitly provided for by a typical trusted operating system. As an example, a typical trusted operating system provides access control to file system objects from user subjects. The level of access control required for the TGG consists of regulating access to destination host objects from source host subjects. This is a level of subject/object granularity inconsistent with typical trusted operating system features. Flow limitation is tightly coupled with access control rules and is similarly inconsistent within an operating system context.

Due to the same subject/object granularity mismatch, typical trusted operating system audit records would not indicate source and destination hosts as the subjects and objects respectively. Audit records may, however, be able to indicate the appropriate subject and object regarding remote management actions. This is dependent on the level of authentication associated with the remote management action. Authentication at a source

level presents a problem, whereas source and user authentication (utilizing typical trusted operating system authentication features) would allow the capture of the pertinent audit information.

Label information may be obtained via a table maintained by the IP software or via the utilization of system calls initiated by the IP software to obtain the level of the single-level channel utilized for communications. The single-level label associated with the communication channel from the low side network would likely be provided via the utilization of operating system calls.

3.2.1.3 Trusted Operating System Performance Calculation

Accurate performance estimation requires a full knowledge of the hardware upon which the gateway functionality is to be performed, the design of the software, and a careful allocation of the available resources to the functionality. In particular, one must pay careful attention to such things as memory cycle time, instruction counts and timings, and the structure of the gateway. A significant effort is required to perform this analysis in detail, and even then, estimates that turn out to be within 20% of the actual results are considered to be very good. For our purposes we devised a simple formula which would give us rough figures for purposes of comparison among trusted operating system products.

A rule of thumb that we obtained from Proteon provided us with the basis for this formula, which states that if I/O tasks have been off loaded to an I/O co-processor, it requires about 500 Motorola 68000 instructions to perform gateway functions for a single one thousand bit packet. Further, the processing time is largely independent of the packet size. The Motorola 68000 is a Complex Instruction Set Computer (CISC) processor. For purposes of comparison only, we extended this rule to all processors (using 1000 instructions per datagram in the case of Reduced Instruction Set Computer (RISC) based CPUs), and then made our performance estimate equal to the minimum of this estimate and the datagram bandwidth of the I/O co-processor, if known.

The performance estimates generated using this approach do not take into account the overhead resulting from the use of a trusted operating system. Such overhead is a consequence of the added checking that must take place and added complexity of context switching. While such penalties have been large in early trusted operating system efforts, current design practices should be able to limit the penalty to around 10% for most cases. In general, the estimates in this section are optimistic. The range, however, is sufficiently above the target range for the TGG so as to provide confidence that these options could satisfy TGG performance requirements.

3.2.2 Trusted Operating System Survey

This section presents the results obtained from the survey of the various trusted operating system products, available or soon to be available, that might be used as a development base for the TGG. For each trusted operating system product a detailed description of the operating system and hardware base is provided followed by an assessment of the products' ability to meet TGG criteria.

CRITERIA SYSTEMS	COMMUNICATION PROTOCOLS				REMOTE MANAGEMENT PROTOCOL		CERTIFICATION		COST
	X25	DDN X25	TCP/ IP	IP/ LABELS	BASES	SECURE BASES	LEVEL	DATE	
AT&T SYSTEM VMLS		X	X	F	X	F	B1	2000	30-40K
BINOS	X		X	F		X	C2/B2		40-50K
GEMINI GEM505									
CUSTOM PROCESSOR	X					X	B3/A1	4000	100K
PC - AT		X	X			X	B3/A1	4000	20-30K
GOULD UTX/328	X	X	X		X		B1	3000	65K
HONEYWELL STOP	F	F	F		X	F	B3	2000	140K
HONEYWELL SECURE UNIX	X	X	X		X		B1	4000	15-60K
IBM SECURE XENIX	X	X			X		B2	2000	12-20K
SUNOS MLS	X	X	X			X	B1	3000	25-45K

F INDICATES TO BE PROVIDED IN A FUTURE VERSION OR PRODUCT

Figure 3-2 Trusted Operating System Capabilities/TGG Criteria Mapping

Figure 3-2, Trusted Operating System Capabilities/TGG Criteria Mapping, provides a summary of the capabilities of the various trusted operating systems with regard to the TGG criteria. This figure represents only those TGG criteria that vary among trusted operating systems. Referring to Figure 3-2, all columns are self explanatory with the possible exception of the remote management column. The two entries under this column equate to the support of the basic fundamentals necessary for evolution to a remote management capability and the support of these basic fundamentals in a secure manner. The later is typically provided only within networks consisting entirely of the particular vendor's products.

With regards to UNIX vendors, throughout the vendor survey their support of the evolving POSIX standard has been indicated. Of particular concern, is the support of a future POSIX secure networking standard. Currently, the POSIX standard does not address networking concerns let alone secure networking. However, the current POSIX standard will be evolved to include secure networking services. Per a recent conversation with Dennis

Steinauer, Chairman of the P1003.6 POSIX security working group, a networking group has only recently been formed at the last POSIX gathering. The networking group will be working in conjunction with the POSIX security working group for the achievement of a POSIX secure networking standard.

At the conclusion of each product survey, an assessment is provided that serves to identify that particular product's suitability for a role such as the TGG. This assessment is in consideration of trusted operating system products only. A subsequent section serves to provide this assessment in light of both trusted operating systems and gateway products. In order to provide this assessment, key criteria must be identified that drive the assessment. Several of the TGG criteria are either consistently not met or are consistently satisfied regardless of the specific trusted operating system involved. The trusted operating system weaknesses associated with the gateway operation protocol and gateway services criteria are applicable to all products as defined in Section 3.2.1, Critical Criteria. On the other hand, all trusted operating system products embody more than adequate performance figures. As a result, the assessment section provided with each product survey will focus on the communication protocols, remote management, cost, and schedule in providing the assessment.

The trusted operating system survey involves the following products:

- AT&T System V/MLS;
- BiiN/OS;
- DEC SEVMS;
- Gemini GEMSOS;
- Gould UTX/32S;
- Honeywell STOP;
- Honeywell Secure UNIX;
- IBM Secure XENIX; and
- SunOS MLS.

3.2.2.1 AT & T System V/MLS

AT&T SYSTEM V/MLS

- **ADVANTAGES**
 - SOON TO BE B1 CERTIFIED
 - DDN X.25 SUPPORT
 - TCP/IP SUPPORT
 - LIMITED SECURITY ENHANCED NETWORKING
- **DISADVANTAGES**
 - COST
- **FUTURE DIRECTIONS**
 - NETWORKING
 - TCP/IP WITH IP SECURITY LABELS
 - REMOTE FILE SHARING WITH LABELS

The AT&T System V/MLS UNIX™ is a multi-level secure version of the AT&T System V UNIX operating system. It is currently targeted at the B1 certification level. The AT&T System V/MLS runs on AT&T's 3B2/600 and 3B4000 computers. The 3B2/600 has been selected as the basis for the TGG assessment.

3.2.2.1.1 Operating System

The System V/MLS operating system supports access control, user identification and authentication, and a limited security-enhanced network capability. To support B1-level requirements some of the standard UNIX features have either been modified or excluded from the operating system. In addition to security, the design goals for System V/MLS include requirements for minimum impact on System V internals, and no impact on the System V Interface Definition. As a result, System V/MLS retains a high degree of compatibility with normal System V UNIX.

Mandatory access control and data labeling in System V/MLS were implemented via an enhancement of the UNIX group feature. All objects were already labeled with a GID (Group ID), and thus security labeling was

implemented by associating a security label with each GID. The kernel was then modified to implement the mandatory protection policy. Modifications of commands controlling groups and GIDs were required to support this approach. Discretionary access control is provided entirely via the standard UNIX owner, group, and other access permissions. Access control lists are not supported.

User identification and authentication has been strengthened via removal of the encrypted passwords from the password/group files, and imposition of limits on root, login, passwd, and super users. For instance, one may no longer log in as root, super users may not obtain root privileges except at the system console (optional), and automatic password generation is supported (optional).

From a secure networking aspect, System V/MLS currently supports a remote file sharing capability that maintains mandatory and discretionary access control labels and identifiers across a System V/MLS network. This functionality is not currently part of the package submitted for evaluation, but plans are to eventually submit this capability for evaluation.

Several modifications, additions, and deletions have been implemented with System V/MLS. Additions include the support of security audit files and device labeling. Modifications have been made to mail, print queues, and temporary files to close covert channels and maintain access control. The deletion of UUCP (UNIX system to UNIX system Copy function), and the imposition of restrictions on the stat (get file STATus) and ps (report Process Status) commands were required to comply with access restrictions and close covert channels.

AT&T System V/MLS is now in formal evaluation for B1 certification. B1 certification is tentatively expected in 2Q89 (refers to calendar year, along with all subsequent uses of this notation). Further modifications to meet B2 and B3 certification requirements are planned, with completion in 2-3 and 3-5 years respectively. Future releases of System V/MLS, specifically Release 4 (due in 3Q89), are expected to comply with all available POSIX standards, including security standards when defined.

3.2.2.1.2 Hardware

System V/MLS runs on a number of AT&T machines. Of these, AT&T's 3B2/600 super microcomputer has been chosen for the TGG application. The 3B2/600 supports multiple processors with cache memory and high capacity main and secondary memory.

The 3B2/600 is based on the WE 32100 CPU which is rated at 2.6 to 4 MIPS. The actual performance figure is dependent on whether one or two CPUs are utilized. A cache memory of 6 KB is available to improve overall data and instruction retrieval times.

Main memory supports ECC and may be configured with four to sixteen MB of RAM. Disk storage of up to 6.5 GB may be provided, however secondary storage in excess of 294 MB requires an expansion cabinet.

Unfortunately, we have been unable to obtain hard data on the availability of a separate I/O processor.

3.2.2.1.3 TGG Criteria Mapping

The following sections map the capabilities of AT&T System V/MLS running on an AT&T 3B2/600 against the TGG mapping criteria.

3.2.2.1.3.1 Communication Protocols

System V/MLS on the 3B2/600 supports a DDN X.25 suite and TCP/IP protocols, the latter via a third party. While the TCP/IP does not currently support security labeling, support for this feature is planned in the future.

3.2.2.1.3.2 Secure Remote Monitoring and Control

Standard UNIX networking capabilities, in the form of UUCP, have been removed from the MLS product. However, System V/MLS does support a remote file sharing capability that maintains mandatory access labels and discretionary access identifiers across a System V/MLS network. This capability is not currently part of the base product submitted for B1 evaluation.

3.2.2.1.3.3 Performance

The 3B2/600 runs at either 2.6 or 4.0 MIPS, depending on whether the optional multi-processor is installed. Assuming that adjunct communications processors can relieve the CPU(s) of the I/O burden, an estimated maximum throughput of 5,200 or 8,000 datagrams per second is achievable. Adjunct communications processor throughput limitations will likely reduce this value. As indicated earlier, we have been unable to obtain data on the availability of an I/O processor and its associated throughput.

3.2.2.1.3.4 Cost

System V/MLS running on a 3B2/600 would cost approximately \$30K to \$40K.

3.2.2.1.3.5 Schedule

AT&T System V/MLS completed developmental evaluation for B1 certification in 3Q88, and commenced formal evaluation in 4Q88. Final B1 certification is expected 2Q89.

3.2.2.1.4 Evaluation

The AT&T System V/MLS - 3B2/600 is a viable candidate for the baseline TGG application. Strong points include the support of the DDN X.25 suite and TCP/IP, a partial secure networking capability, more than adequate performance, and a timely B1 certification schedule. The negatives are cost, which is the middle of the range for trusted operating system products, and uncertainty over the availability of a communications co-processor.

3.2.2.2 BiiN/OS

BiiN 20 SYSTEM

• ADVANTAGES

- TCP/IP SUPPORT
- HIGH PERFORMANCE
- SECURITY ENHANCED NETWORKING

• DISADVANTAGES

- COMMERCIAL X.25 SUPPORT ONLY
- COST
- B-LEVEL EVALUATION TIME FRAME UNSPECIFIED

• FUTURE DIRECTIONS

- B-LEVEL COMMITMENT
- NETWORKING SECURITY COMMITMENT

BiiN is a joint venture of Intel and Siemens. It was created to develop a line of general purpose, multi-processor computers directed at the critical applications market. At present, this line contains two machines: the BiiN 20, which is available in both one and two processor configurations; and the BiiN 60, which is available in 2, 4, 6, and 8 processor configurations. Our assessment is directed at the single processor version of the BiiN 20.

3.2.2.2.1 BiiN 20 Hardware

Like all BiiN machines, the BiiN 20 is constructed around four custom VLSI components: the CPU, the CP (Channel Processor), the BXU (Bus eXchange Unit), and the MCU (Memory Control Unit). A brief outline of the function and capabilities of each of these components follows:

CPU: The CPU is essentially an Intel 80960 with a number of enhancements. Some of the more notable of these are on chip floating point support, memory management, and hardware implementation of a number of operating system functions including queue management and address range checking for capability based access control. The CPU is a full

32 bit tagged processor, and runs at about 5.5 MIPS and 1.0 MFLOPS (on 32 bit values) depending on cache size and hit rate.

CP: The CP is a specialized processor optimized for I/O control and support. It is capable of sustaining transfer rates of up to 32 MB/second. The CP is a processor in its own right, and thus even the single processor version of the BiiN 20 really has two processors which communicate with each other via shared memory and interrupts. The CP could be programmed to perform communications functions thereby offloading the CPU.

BXU: The BXU provides bus interface, cache coherency and cache management services to a CPU and/or a CP.

MCU: In addition to its memory control functions, the MCU provides its own bus interface, performs ECC functions on both address and data lines, and supports memory scrubbing. The MCU also supports a spare bit for on line replacement of a failed DRAM chip.

Each of these components have extensive built in self check capabilities.

The single processor version of the BiiN 20 consists of a memory board, and a processor & I/O board connected to a 40 MB/sec system bus. The processor & I/O board contains a CPU and a CP sharing a BXU and a cache. The memory board consists of an MCU and either 8 or 16 MB of RAM. The base version of the BiiN 20 comes with a 180 MB hard disk and a tape drive. Disk storage capacity is expandable to 1.7 GB per channel. The BiiN 20 performs extensive self checks upon startup.

3.2.2.2 BiiN Operating System

The BiiN operating system, known as BiiN/OS, is unusual in that many of its primitive functions are implemented in hardware. For instance, the context switch required to respond to an interrupt is handled completely in hardware. As a result of this design feature, BiiN machines can respond very quickly to external events, and seem well suited to real time applications. A more important example of this approach is the hardware support for BiiN/OS's capability based addressing and protection scheme.

The virtual memory in any BiiN machine is divided up into objects ranging in size from 2^6 to 2^{32} bytes. Objects are accessed via access descriptors, which may be thought of as combination pointers and access rights specifiers (i.e. read, write, execute, etc.). While the user of a BiiN machine sees a 32 bit word, the actual word size is 33 bits. The 33rd bit is not accessible to unprivileged processes, and is used to flag object descriptors. Combined with appropriate operating system design, this hardware feature prevents corruption of object descriptors, and thereby prevents unauthorized access to

objects. In addition to providing the foundation of BiiN/OS's security features, this arrangement also permits hardware enforcement of the object oriented design principles that are used throughout the software provided with the BiiN machines.

From the user perspective, security under BiiN/OS is based on discretionary access control lists. BiiN machines are intended to function in networks, and thus discretionary access control is maintained across BiiN homogeneous networks. BiiN's initial certification target is C2, with an ultimate goal of B2 or better.

BiiN/OS also includes a complete application development environment, designed with particular attention to the security aspects of the process. In addition, BiiN/OS offers a operating system primitive level UNIX System V interface, which is committed to conform to the POSIX standard, and a UNIX style software development environment.

Other notable features of BiiN/OS include transparent distributed processing across a network, transaction processing services, ability to maintain an audit trail on all changes to a specified file, and compatibility with numerous communications and network protocols.

3.2.2.2.3 Upgrade Options

As indicated earlier, the single processor BiiN 20 is the bottom of BiiN's product line, and as such lacks some of the features of the BiiN 60. Software developed for the single processor BiiN 20 is binary compatible with the other members of the BiiN family. While the BiiN 20 and the BiiN 60 share many components, they are two different machines, and BiiN 20 hardware cannot be upgraded to a BiiN 60.

The dual processor BiiN 20 and all configurations of the BiiN 60 are multi-processor machines with a shared memory on a bus architecture. A large amount of the multi-processor scheduling problem is handled in hardware. Given the nature of the proposed applications, expected performance would improve more or less linearly with the number of CPUs and CPs within the limitations of bus bandwidth and cache size. The BiiN 20 uses a single 40 MB/sec bus, the BiiN 60 has two such busses which operate in parallel. The BiiN 60 can continue to operate in a degraded mode if one bus fails.

In addition to the features of the BiiN 20, the major selling points of the BiiN 60 are reliability, fault tolerance and increased processing power. For reliability, the BiiN 60 is offered in configurations with duplicates of most essential components, and automatic fault detection and failover capabilities. Failed components are reported to the main console for replacement. In most cases, the failed component can be replaced without bringing the machine down. The BiiN 60's increased processing capability comes from its ability to support up to eight CPUs. Fault tolerance is obtained via shadowing of

components, and is transparent to applications programs. In the case of CPUs, three levels of fault tolerance are available: none, processor shadowing for fault detection with automatic re-attempt in the event that a fault is discovered, or quad redundancy for continuous service in the event of a hardware failure.

3.2.2.2.4 TGG Criteria Mapping

The following sections map the capabilities of the single processor BiiN 20 against the TGG mapping criteria.

3.2.2.2.4.1 Communication Protocols

The BiiN 20 supports the X.25 and TCP/IP protocols, although the X.25 is not a DDN certified version. The TCP/IP protocol does not support security labeling, however the vendor has indicated a willingness to add this feature if required.

3.2.2.2.4.2 Secure Remote Monitoring and Control

The BiiN machines maintain access control across a BiiN homogeneous network, and thereby provides a mechanism for secure remote monitoring and control. Unfortunately, access control is not maintained across a heterogeneous network. The current version of the BiiN/OS supports discretionary access control only.

3.2.2.3.4.3 Performance

The BiiN CPU runs at 5.5 MIPS, and the CP (Channel Processor) is capable of supporting aggregate transfer rates through the BiiN 20's serial links of up to 5 Mbits/sec. Based on the quoted performance figures, an estimated throughput of 5,000 datagrams/second could be supported. However, the BiiN 20 is limited to six serial ports, and thus line speed limitations could be a source of difficulty.

3.2.2.2.4.4 Cost

Depending on configuration, BiiN 20s range in price from \$40K to \$80K. In all probability, a minimal configuration would suffice for the TGG role.

3.2.2.2.4.5 Schedule

The BiiN product line was unveiled early in 4Q88. The BiiN 60 was available in 4Q88, and the BiiN 20 will become available in 1Q89. BiiN is now involved in a C2 developmental evaluation. Plans for a B-level certification are unavailable. BiiN has made an internal commitment to support full B2 functionality by the end of 1990.

3.2.2.5 Evaluation

The lack of a B-level certified product for the BiiN/OS is a very serious disadvantage that may not be resolvable within the TGG implementation timeframe. The BiiN/OS has just begun the evaluation process at the C2 level. Scheduling for a B-level product is not predictable at this time. Cost may also prove to be a detriment depending on the specific hardware configuration required. Costs vary from the middle to the high end of the operating system range. Otherwise, the BiiN/OS provides other capabilities required for the TGG including the support of TCP/IP, security enhanced networking, and high performance.

3.2.2.3 DEC SEVMS

DEC SEVMS

- **ADVANTAGES**

INSUFFICIENT DATA
TO COMPLETE
EVALUATION

- **DISADVANTAGES**

DEC SEVMS (Security Enhanced VMS) is a security enhancement package intended to upgrade the level of security offered by the VAX/VMS operating system. VAX/VMS already has an extensive set of optional security features, and thus the main thrust of SEVMS is to add mandatory access controls and labeling to the VMS operating system. The current version of SEVMS is C2 certified, with a targeted certification at the B1-level. The MicroVax II appears to be a suitable base for SEVMS in the TGG application, however we have been unable to obtain sufficient information for a proper assessment of this solution to the TGG problem.

Vendor contact to date has consisted of several phone conversations. The initial conversation was utilized to introduce the TGG and its requirements, and to request product literature. Several subsequent phone

conversations were held in an attempt to set up a technical meeting with DEC to discuss the TGG and DEC products in more detail. Attempts to set up a meeting have been unsuccessful.

3.2.2.3.1 Operating System

As indicated earlier, unaugmented VAX/VMS offers an extensive set of security features, and thus many of the features provided by SEVMS are provided in some form by unaugmented VAX/VMS as well. For example, both operating systems support several varieties of password protection for system access control, and in both cases discretionary file access control may be based on either the System/Owner/Group/World model, on the access control list approach, or on a combination of the two. Thus the following discussion of SEVMS mentions many features which will be familiar to anyone who is familiar with unaugmented VAX/VMS.

SEVMS supports several forms of password protection. These features include password generation, length and lifetime parameters, and various levels of required passwords. Users may use a built-in password generator in the selection of their passwords. At the system administrator's option, the use of the password generator may be made mandatory. The minimum length and lifetime of passwords may also be enforced as indicated in the user's account definition. Accounts of a sensitive nature, such as a system administrator, may require two levels of passwords in order to login. Before an account of this type can become active, a primary followed by a secondary password is required.

Discretionary access control is based on a User Identification Code (UIC). The UIC is used in conjunction with a System/Owner/Group/World structure access protection scheme that matches the UIC of the subject (e.g., user) against that of the object (e.g., file) to determine access. Within each category, access may be indicated as read, write, execute, delete, and/or control. Control access given to an individual by the object owner gives that individual the ability to set new access permissions for that object. Access control lists are provided to permit fine tuning of the UIC based protection scheme to individual identifiers (e.g., users) as required.

A very flexible auditing capability is provided for the logging and reporting of auditable events. Individual auditable events may be selectively enabled or disabled by an individual with the security privilege. This individual is typically the security administrator. Auditable events that SEVMS supports include login/logout, object access, access rights modifications, classification modifications, and software installations. An audit analysis program is available that provides the security administrator with a flexible means of generating audit reports.

SEVMS is available now as a C2 certified product. A B1 version, which includes mandatory access control, is currently undergoing evaluation.

3.2.2.3.2 Hardware

As indicated earlier, it appears that the MicroVax II would be an appropriate hardware base for SEVMS running the TGG application. We have been unable to obtain sufficient data from DEC about the MicroVax II to complete this section. Indeed, much of the information in the previous section was drawn from our staff's prior knowledge of the VMS and SEVMS operating systems.

3.2.2.3.3 TGG Criteria Mapping

Due to insufficient data, we are unable to map the capabilities of SEVMS running on a MicroVax II against the TGG mapping criteria.

3.2.2.3.4 Evaluation

We lack sufficient data for an evaluation.

3.2.2.4 Gemini GEMSOS

GEMINI GEMSOS

- **ADVANTAGES**
 - SOON TO BE A1 CERTIFIED
 - SECURITY ENHANCED NETWORKING
 - ADEQUATE PERFORMANCE
 - COST (FOR PC BASED)
- **DISADVANTAGES**
 - BASIC X.25 SUPPORT ONLY
 - SECURITY ONLY FOR HOMOGENEOUS NETWORK
 - COST (FOR GEMINI PROPRIETARY HW)
- **FUTURE DIRECTIONS**
 - DEVELOPMENT OF A STANDARD UNIX INTERFACE

Gemini GEMSOS is a multi-level secure, multiprocessing operating system directed at the B3 level of certification. When running on Gemini's

proprietary 80286/80386 based hardware, GEMSOS is capable of controlling multiple processors. GEMSOS is also available on a modified IBM PC AT or compatible.

3.2.2.4.1 Operating System

The GEMSOS operating system provides mandatory security and integrity policy enforcement, secure-block oriented secondary storage I/O, secure device (i.e. terminal, printer, etc.) I/O, application-directed segment-based primary memory management services, multiprocessing primitives for process synchronization and inter-process communication, network security and, on Gemini's own hardware, transparent management of up to eight tightly coupled processors.

Mandatory access control rules in GEMSOS are enforced by a security kernel. All entities (i.e. processes, secondary storage volumes and segments, devices, etc.) possess an access class. This access class specifies not only the sensitivity class of the entity, but its integrity class as well. The security kernel uses the access classes of entities to enforce the *-property and simple security condition.

GEMSOS does not support a file system. Instead, each volume of secondary store is divided into segments of 0 to 64 Kb. These segments may be swapped into and out of central memory at software request. Note that segments and volumes are subject to the same security constraints as any other entity in GEMSOS.

GEMSOS supports multi-process synchronization via event counts and sequencers. When running on a multi-processor Gemini machine, GEMSOS also supports multiple processors in a fashion that is transparent to the user.

Gemini offers a package of trusted software to run under GEMSOS. Of particular interest with regard to the TGG is software supporting transaction based, multi-level secure communications between machines running GEMSOS. Multi-level secure loosely coupled networks are also supported. In this case as well, security is only maintained in GEMSOS homogeneous networks.

3.2.2.4.2 Hardware

GEMSOS is built around the Intel 80286, and relies heavily on the security features of this CPU. GEMSOS is available on two machines: the IBM AT (or selected AT compatible), and the Gemini TCB (Trusted Computer Base) machine.

Gemini's TCB is a Multibus based machine which is available in many configurations. In its largest configuration, it can support up to eight Intel 80286 or 80386 CPUs operating concurrently. Communication between processors is implemented via shared memory. Global memory is at least 0.5 MB, and processor local memory is at least one MB. Upgrades are available.

Hardware support for the NBS DES is provided. These machines are available with no secondary storage, with floppy disk drives, with hard disk drives, and/or with tape drives. Disk storage upwards of 140 MB is available. In addition, a number of network interface cards are available such as an ETHERNET LAN and X.25 controllers.

GEMSOS is also available in the form of a MLS IBM AT workstation upgrade kit. This kit includes NBS DES encryption hardware and a special GEMSOS PROM.

3.2.2.4.3 TGG Criteria Mapping

The following sections map the capabilities of GEMSOS running on either of its two hardware bases against the TGG mapping criteria.

3.2.2.4.3.1 Communication Protocols

GEMSOS on the Gemini TCB supports basic X.25 via a network interface card. At present, this card requires a dedicated CPU for its management, however a new X.25 interface is expected in the next three to six months which may alleviate this requirement. Gemini offers no support for TCP/IP on its TCB machines.

On the IBM AT, standard AT bus expansion cards may be used, however their use requires the construction of trusted hardware drivers. Gemini has indicated a willingness to support such an effort. Thus, I/O coprocessor cards supporting X.25 and TCP/IP which are available from third parties (i.e. Frontier Technologies) could be used to provide the necessary communications services.

3.2.2.4.3.2 Secure Remote Monitoring and Control

Gemini provides secure network management functionality as part of the non-kernel TCB package. The secure networking capability allows the creation of remote processes with the same attributes as the creating process. This includes both mandatory and discretionary access control attributes. The secure networking capability is available only between Gemini equipment on a homogeneous network.

3.2.2.4.3.3 Performance

Depending on the hardware selected, machines running GEMSOS run at 0.5 to 24 MIPS. Assuming that I/O coprocessors can relieve the CPU(s) of the I/O burden, an estimated maximum throughput of 1,000 to 48,000 datagrams per second performance would be obtainable. With the exception of slow ATs, I/O coprocessor throughput limitations will doubtless be the limiting factor. In the case of the AT, currently available I/O co-processor boards are capable of handling only a fraction of this load. However the TGG throughput requirements could be met through the use of multiple boards.

3.2.2.4.3.4 Cost

Depending on the configuration selected, GEMSOS systems run from \$20K on up into the hundreds of thousands of dollars. If the IBM AT option is selected, the cost should be in the \$20K to \$30K range.

3.2.2.4.3.5 Schedule

GEMSOS and its supporting hardware are currently available. B3 formal evaluation commenced in 3Q88, and is expected to be complete in 4Q89.

3.2.2.4.4 Evaluation

Cost may prove to be a major factor in the consideration of the majority of Gemini products as a base for the TGG. The PC AT or compatible version with the Gemini security upgrade kit represents the most cost effective solution. Aside from the cost factor, Gemini products provide security enhanced networking, adequate performance, and a timely certification schedule.

3.2.2.5 Gould UTX/32S

GOULD UTX-32S SECURE UNIX

- **ADVANTAGES**
 - SOON TO BE B1 CERTIFIED
 - DDN X.25 SUPPORTED
 - TCP/IP SUPPORTED
 - ADEQUATE PERFORMANCE
- **DISADVANTAGES**
 - COST

The UTX/32S operating system is a secure operating system based on Gould's Universal Time-Sharing Executive (UTX/32) operating system, which is derived from Berkeley BSD 4.X and Bell Laboratories System V UNIX. The current version is C2 certified. The next version of UTX/32S is targeted for the B1 level with certification expected in either 2Q89 or 3Q89. UTX/32S runs on the Gould PowerNode super minicomputer series. The PowerNode 6000 has been identified as the most likely candidate for the TGG application. This machine is a 32-bit multiple processor capable minicomputer.

3.2.2.5.1 Operating System

In support of C2-class requirements, UTX/32S supports access control, user identification and authentication, accountability, and secure inter-process communication. In order to provide a secure UTX/32S operating system, some features of the baseline UTX/32 operating system have been deleted. To support B1-class requirements a subsequent version of UTX/32S, currently undergoing formal evaluation, will support mandatory access control and associated labeling.

Discretionary access control mechanisms utilize restricted environments and a modified group mechanism. These mechanisms are in addition to the normal UNIX owner, group, and other file access permissions that are available. The modified group mechanism limits user activity to one group at a time. The utilization of restricted environments places all users into one of two user high-level categories: unprivileged or privileged. Administrative users belong to the privileged category and utilize the associated privileges in order to perform the required administrative duties. All other users are considered as unprivileged and are placed in a restricted environment, of which there may be many dependent on the functional partitions of the specific system. A restricted environment exists as a sub-tree of the UNIX file system into which a user is placed upon session startup. The restricted environment is viewed as a discrete UNIX file system subject to the owner, group, and other file permissions as defined. As a result, unprivileged users cannot access information outside of the restricted environment, making it possible to keep system-sensitive files protected. Mandatory access control will be provided with the B1 version of UTX-32S.

UTX/32S supports a user ID and logging ID versus the standard UNIX real user ID and effective user ID. The logging ID is the permanent identification of the user required to log on and never changes during the session. This is the ID that is used for individual accountability functions such as auditing. The user ID reflects the user's current privileges. As an example, the user ID would reflect the user's current ID as a result of a setuser system call. The incorporation of the logging ID allows UTX/32S to provide an unambiguous accountability of all individual user actions.

UTX/32S provides for the generation of audit data pertaining to critical events and tools for the subsequent analysis and reporting of those critical events. Auditing of critical events may be categorized as actions against files, directories, and processes.

UTX/32S provides for secure inter-process communication via secure sockets to the system's trusted servers. Secure inter-process communication is required due to the nature of the information passed to/from the trusted servers. Trusted servers consist of services that perform user authorization information processing, mail, printing, device control, and administrative processing.

In order to provide for a secure UTX/32S operating system, several standard features of UNIX have either been deleted or modified. System calls that set or return real or effective user and group identification have been either restricted or deleted. Additionally, the set user and group identification bits have been deleted from the file permissions, and several system calls dealing with these bits have been modified. These include system calls that determine accessibility, change access permissions, and retrieve user or group identification values.

Gould is committed to supporting the POSIX standards in UTX/32S as the standards become available.

3.2.2.5.2 Hardware

The PowerNode 6000 32-bit super minicomputer lies at the low-end of Gould's PowerNode super minicomputer series, and has been targeted at applications such as data communication gateway or file server. The PowerNode 6000 supports multiple processors with virtual memory addressing and cache memory; separate input/output processor(s); and high-capacity main and secondary memory.

The PowerNode 6000 has the ability to support multiple processors with a performance figure of 1 to 3 MIPS dependent on CPU configuration. In addition to the standard CPU that performs input/output processing, interrupt processing, and computational tasks an additional duplicate Internal Processing Unit (IPU) is available. The IPU serves to off-load the computational tasks. Both the CPU and IPU support virtual memory addressing up to 16 MB and 32 KB of cache memory. The CPU, IPU, memory modules, and input/output processor all reside on the SelBUS. The SelBUS is a high-speed synchronous bus with a bandwidth of 26.67 MB per second.

A separate input/output processor is available that operates independently of and in parallel with the CPU, provides a medium level of performance, and supports several device controllers. The input/output processor is capable of providing data transfers at a rate up to 1.5 MB per second. A single input/output processor can support up to 16 device controllers exclusive of high-speed devices such as disks or tapes which interface directly to the SelBUS.

Main memory, which may be configured in 1 MB increments from 2MB up to 16 MB, is implemented with high-speed dynamic RAM with error correcting code. Main memory may be interleaved to achieve a higher level of performance. A memory protection scheme is implemented that requires privileged operation to alter memory protection registers. Disk storage in excess of 2 GB is supported.

The PowerNode 6000 is also available in a single board version, which is part of Gould's SelCONNECTION product line. This version carries approximately the same price tag as the multi-board implementation.

3.2.2.5.3 TGG Criteria Mapping

The following sections map the capabilities of UTX/32S running on a PowerNode 6000 against the TGG criteria.

3.2.2.5.3.1 Communication Protocols

The UTX/32S - PowerNode 6000 solution supports both the DDN X.25 suite and TCP/IP protocols. The DDN X.25 suite is supported by both hardware and software. Level 2 of the X.25 service is supported by Gould's Synchronous Communications Multiplexer, whereas Level 3 is supported by software. TCP/IP, with IP security labeling capability (IP option 133), is supported by the PowerComm DDN Interface Module. In addition to TCP/IP, this software supports FTP, TELNET, and SMTP.

3.2.2.5.3.2 Secure Remote Monitoring and Control

UTX/32S does support network interaction via UNIX cu (call UNIX) and uucp (UNIX-to-UNIX copy) to appropriately equipped UNIX-based computers. This is a basic foundation for remote monitoring and control, however it is not via a trusted mechanism. Currently, the UTX/32S product does not support any sort of secure network functionality that could serve as a basis for secure remote monitoring and control functionality. Specific future directions in this area are not decided at this time; however, Gould is committed to the evolving POSIX standard.

3.2.2.5.3.3 Performance

Based on various CPU configurations providing 1 - 3 MIPS of performance and a dedicated I/O processor capable of 1.5 MB per second transfers, datagram processing performance is upwards of 1,000 datagrams per second.

3.2.2.5.3.4 Cost

The base product costs approximately \$65K.

3.2.2.5.3.5 Schedule

UTX/32S is currently available as a C2 certified secure operating system. Beta release for the B1 product is to occur 4Q88 followed by the commencement of B1 formal evaluation 1Q89. Final B1 certification is expected 2Q89 or 3Q89.

3.2.2.5.4 Evaluation

The cost of the UTX/32S - PowerNode 6000 is at the high end of the operating system range, which is a negative factor in considering this option for the TGG. The factor may be adjusted somewhat by discounts that may be available. Aside from cost factors, the Gould solution provides support for

the DDN X.25 and TCP/IP protocols along with IP labeling, has good performance, and has a timely B1 certification schedule.

3.2.2.6 Honeywell SCOMP Trusted Operating Program

HONEYWELL STOP

- **ADVANTAGES**
 - SOON TO BE B3 CERTIFIED
- **DISADVANTAGES**
 - COST
 - LACK OF X.25 SUPPORT

The STOP operating system is an A1 certified secure operating system originally developed by Honeywell for the Secure Communications Processor (SCOMP). The SCOMP is based on a Honeywell DPS 6 16-bit minicomputer with a modified CPU and a Security Protection Module. Realizing the limitations of a 16-bit minicomputer, Honeywell has ported the STOP operating system to a DPS 6 PLUS 32-bit minicomputer, and named the resulting system the Honeywell TX-200. This port (STOP version 3.0) is initially directed at the B3 level, with certification expected in 2Q89.

3.2.2.6.1 Operating System

The Honeywell STOP version 3.0 is a multi-level secure operating system which supports multiple secrecy and need to know categories. The STOP security kernel enforces both mandatory secrecy and integrity policies. Access control lists are also supported. STOP includes trusted software to perform such global services as I/O management, inter-process communication, network services, and file system management. STOP supports extensive auditing capabilities.

STOP version 3.0's application programming environment is provided by the Commodity Application System Services (CASS). In order to provide a wide range of applications capability in the trusted environment the user interface has been developed to be similar to that provided by the UNIX System V Interface Definition, with the exception of those services which violate the security policy. This allows the system to support a wide selection of existing applications and reduces the cost of software development.

3.2.2.6.2 Hardware

The Honeywell DPS 6 PLUS 400 Series lies at the low-end of Honeywell's DPS 6 PLUS family. These machines support one to four custom designed 32 bit VLSI CPUs which share the 16 KB cache and main memory. In a multiple processor configuration the CPUs operate in a peer-to-peer relationship sharing the processing load. Virtual addressing, communications co-processor(s), and high capacity main and secondary memory are supported. Virtual addressing is provided by a custom-designed VLSI Virtual Memory Management Unit which provides 2 GB of virtual memory space per process.

Main memory is expandable in 2 MB and 4 MB increments for a total of 16 MB of dynamic RAM. Memory expansion beyond 8 MB requires the inclusion of an expansion cabinet. Memory interleaving is supported, but requires a configuration that supports an additional memory controller configured in an expansion cabinet. Disk storage of 5 GB is supported. A maximum of 3.3 GB is supported without the necessity of an expansion cabinet.

Software disk-caching is supported. This permits many disk I/O functions to be performed at solid state memory access speeds.

A Secure Communications Subsystem has been developed for use with the DPS 6 PLUS running the STOP operating system. The purpose of this module is to permit the use of commercially available VME based communications boards with the TX-200. Due to the need to construct trusted hardware drivers for each such communications board, this module is currently available only for an Ethernet board with TCP/IP. However the construction of the necessary trusted drivers for an X.25 board is in progress, with a projected release in 1Q90. The throughput capabilities of the X.25 board are unknown at this time, however at the very least, they should be sufficient to drive a 56 KB/sec serial line.

3.2.2.6.3 TGG Criteria Mapping

The following sections map the capabilities of STOP operating system running on a DPS 6 PLUS against the TGG criteria.

3.2.2.6.3.1 Communication Protocols

While the TX-200 solution does not now support a DDN X.25 co-processor board, it should do so by 1Q90. TCP/IP is currently supported as part of an Ethernet board & software package, TCP/IP (with option 133) will be available with the X.25 board.

3.2.2.6.3.2 Secure Remote Monitoring and Control

A secure networking capability, which forms the basis for a secure remote monitoring and control capability, is not supported by the STOP operating system. However the vendor has indicated their intention to

support secure networking as soon as standards for same are developed. The TX-200 console can be assigned to any link, and thus Remote monitoring and control could be effected via login over any suitably secure communications link.

3.2.2.6.3.3 Performance

The DPS 6 PLUS can support one to four CPUs for a processing power of 1 to 3.5 MIPS. Thus our performance estimation rule predicts packet throughput in the order of two to seven thousand 1000 bit packets per second, as limited by the communications co-processor boards. Unfortunately, hard throughput data on the communications boards is not available at this time, so it is hard to estimate what limit they would place on overall throughput. However, the tentative information at our disposal indicates that a TX-200 with four X.25 boards should be capable of meeting the TGG throughput requirements.

3.2.2.6.3.4 Cost

The base product ranges in cost from \$80K to \$350K with an average system cost of approximately \$140K.

3.2.2.6.3.5 Schedule

The STOP operating system is currently undergoing formal evaluation. B3 certification is expected 2Q89.

3.2.2.6.4 Evaluation

The Honeywell STOP - DPS 6 PLUS solution is at the high end of the cost range for even a minimal configuration. Further, it does not now support DDN X.25. On the plus side is a realistic B3 certification schedule and previous experience with the certification process involving the A1 SCOMP - STOP product.

3.2.2.7 Honeywell Secure UNIX

HONEYWELL SECURE UNIX

- ADVANTAGES

- SOON TO BE B1 CERTIFIED
- COST
- DDN X.25 SUPPORTED
- TCP/IP SUPPORTED

- DISADVANTAGES

- NO SECURE NETWORKING

Honeywell's Secure UNIX, herein known as HFSI (Honeywell Federal Systems Inc.) B1 UNIX, is a secure operating system based on System V Release 3 UNIX. The HFSI B1 UNIX operating system design concept has placed an emphasis on securing UNIX, not necessarily producing a system that resembles UNIX. Final certification is expected 4Q89. The HFSI B1 UNIX hardware base is the Honeywell Bull XPS-100 computer series, of which the Model X-22 has been identified for the TGG application.

3.2.2.7.1 Operating System

The HFSI B1 UNIX operating system supports access control, user identification and authorization, accountability, and superuser partitions. Implementation of these features required the modification of some UNIX features (i.e. SUID & SGID), and the deletion of others (i.e. Debug, & UID zero).

Both mandatory and discretionary access control are supported by HFSI B1 UNIX. In support of mandatory access control, labeling is provided that distinguishes the sensitivity level of all subjects and objects including devices. In addition to the standard UNIX owner, group, and other file access permissions, discretionary access control is provided in the form of access control lists (ACL). ACLs provide a finer level of granularity of discretionary access control than the standard UNIX file permissions.

User identification and authorization provides the basis for accountability. Accountability, in the form of auditing, provides for the

logging of system unique events, object events, and subject events. System unique events consist of system-level occurrences such as memory faults. Object events consist of actions taken against objects such as a file open or close operation. Subject events consist of user actions such as logging in or out. The definition of which auditable events are active is selectable. Additionally, audit data reduction and reporting capabilities are provided.

Superuser partitions are provided that segregate the UNIX superuser into various roles. Five roles are currently defined that give and constrain permissions according to the defined role. These roles are: security officer, system administrator, system operator, backup and restoration, and file repair. Roles are validated via: a role login plus password, user login plus password, user identified in ACL for role directory, and login at the proper terminal. All of these steps must be successful in order to establish the desired role. As a result of the superuser partition, UID zero (superuser) is no longer supported.

The HFSI B1 UNIX operating system will support the POSIX standards as they become available.

3.2.2.7.2 Hardware

The Honeywell Bull XPS-100 Model X-22 lies at the low end of the Honeywell Bull XPS-100 Series. The Model X-22 supports virtual memory addressing, cache memory, separate input/output processor, LAN module, and a variety of main and secondary memory configurations.

The Model X-22 is a single CPU processor based on the Motorola 68020 32-bit microprocessor. The CPU is clocked at a 16.67 MHZ rate to provide performance in the range of 1.7 to 2.1 MIPS. The actual performance figure is directly dependent on the inclusion of the optional 16 KB associative cache memory. Virtual memory addressing is supported that provides a 16 MB virtual memory space per process.

A communications co-processor board known as a line processor board is available. It is capable of driving two 56 Kbit/sec lines simultaneously. This board supports DDN X.25.

A Local Area Network (LAN) controller, which supports ETHERNET, is available. The LAN controller supports up to 64 users, includes an Intel 80186 microprocessor, 512 KB dual port memory, a high speed direct memory access channel, and extensive onboard diagnostics. Additionally, TCP/IP firmware is packaged with the LAN controller.

Main memory is expandable in 2 MB, 4 MB, or 8 MB increments for a total of 16 MB. Memory interleaving is not supported. Disk storage in excess of 900 MB is available.

3.2.2.7.3 TGG Criteria Mapping

The following sections map the capabilities of HFSI B1 UNIX running on a Model X-22 against the TGG criteria.

3.2.2.7.3.1 Communication Protocols

DDN X.25 and TCP/IP are supported. The TCP/IP protocol supports security labeling.

3.2.2.7.3.2 Secure Remote Monitoring and Control

At present, HFSI B1 UNIX does not support secure networking, although there are plans for this in the future. Untrusted networking is supported.

3.2.2.7.3.3 Performance

The XPS-100 model X22 runs at about 2 MIPS, so our performance estimation rule predicts a maximum packet throughput in the order of four thousand 1000 bit packets per second, as limited by the communications co-processor boards. Thus the Model X22 equipped with four line processor boards should be able to meet the TGG throughput requirements.

3.2.2.7.3.4 Cost

The base product ranges in price from \$15K to \$60K. A suitable base for the TGG would probably lie in the lower middle part of this range.

3.2.2.7.3.5 Schedule

The HFSI B1 UNIX product is currently in the evaluation process. Certification is expected in 4Q89.

3.2.2.7.4 Evaluation

Other than the lack of secure networking facilities, the HFSI B1 UNIX / XPS-100 Model X-22 solution appears to have no major negatives which are not endemic to the trusted operating system approach. The positive side includes cost, support of DDN X.25 and TCP/IP, adequate performance, and a timely certification schedule.

3.2.2.8 IBM Secure XENIX

IBM SECURE XENIX**• ADVANTAGES**

- SOON TO BE B2 CERTIFIED
- ADEQUATE PERFORMANCE
- COST
- THIRD PARTY DDN X.25 AND TCP/IP

• DISADVANTAGES

- NO REMOTE MANAGEMENT SUPPORT

Secure XENIX is a secure version of the XENIX operating system developed by IBM for the PC/AT and PS/2 Models 50, 60, and 80. Secure XENIX is currently targeted for a B2 level with an expected certification 2Q89. IBM is also developing a Secure Trusted Application-level Gateway (STAG) which is based on the Secure XENIX product. The STAG is intended to provide file transfer capabilities among local nodes with the addition of security checking and encryption support.

3.2.2.8.1 Operating System

Secure XENIX is a multilevel secure operating system with multi-processing capabilities, security policy enforcement, and user identification and authorization. Secure XENIX has completed the B2 developmental evaluation phase. B2 formal evaluation commenced in December of 1987. Completion of the formal evaluation is expected in the second quarter of 1989.

The product uses a secure kernel for resource control, security policy enforcement, and user identification and authentication. Implementation of the secure kernel has forced a number of structural changes from XENIX. Most notable amongst these are the removal of the list of encoded passwords from the etc/passwd file, and the removal of any special privileges from the root ID. In addition, a number of system management and configuration operator commands are no longer supported, and some of those remaining have been modified to avoid conflicts with security policy. Given sufficient disk space, Secure XENIX is capable of maintaining extensive audit files.

Secure XENIX is binary compatible with code developed under IBM Personal Computer XENIX versions 1.0 and 2.0, and the previous version of Secure XENIX. Programs written according to AT&T Bell Laboratories UNIX System V Interface Definition are source code compatible with Secure XENIX provided they use only the functions supported by Secure XENIX. Secure XENIX includes a software development environment and a document preparation package. The software development environment is limited to C and 80286 assembler.

There are no plans for implementing any parts of the POSIX standard in Secure XENIX. Indeed, it appears that there are few plans for further development of the Secure XENIX operating system. Instead, the emphasis seems to be on the development of a secure version of IBM's AIX operating system. We are told that this product will support POSIX.

3.2.2.8.2 Hardware

Secure XENIX is intended to run on either an IBM AT or an IBM PS/2 model 50, 60, or 80 without hardware modifications beyond upgrades required to meet the minimum RAM and disk drive capacity requirements. Machines of this class run an Intel 80286 or 80386 CPU at 6 to 25 MHz, with typically 0.5 to 4.0 MB of RAM, a 20 to 100 MB hard disk drive, and one or two floppy disk drives. These systems usually include a 16 or 32 bit wide expansion bus running at 6 to 12 MHz. I/O is normally handled by the CPU, although there are expansion boards available that can relieve the CPU of some of the I/O burden. An ETHERNET communication adapter is available that is used in conjunction with the TCP/IP protocol.

3.2.2.8.3 TGG Criteria Mapping

The following sections map the capabilities of IBM Secure XENIX against the TGG mapping criteria.

3.2.2.8.3.1 Communication Protocols

Secure XENIX supports TCP/IP (without security labeling) as part of an optional ETHERNET LAN capability. There are third party sources (i.e. Frontier Technology) for communications co-processor boards which support the DDN X.25 & TCP/IP protocols. Use of such boards would require the construction of a suitable trusted hardware driver.

3.2.2.8.3.2 Secure Remote Monitoring and Control

Secure XENIX does not offer a trusted mechanism for this function, although it does offer a remote login and file transfer facility which could provide a basis for remote monitoring and control.

3.2.2.8.3.3 Performance

Depending on the AT or PS/2 selected, the 80286 or 80386 (in 80286 mode) running under Secure XENIX would be capable of at most 4.5 MIPS. Dedicated communications I/O co-processor expansion boards could be used to relieve the CPU of the basic network I/O tasks. Given this scenario, an estimated maximum throughput of 9000 one thousand bit datagrams per second could be achieved. Unfortunately, currently available I/O co-processor boards are capable of handling only a fraction of this load. However the TGG throughput requirements could be met through the use of multiple boards.

3.2.2.8.3.4 Cost

With I/O co-processor boards and Secure XENIX, the price of a suitable platform would lie in the range of \$12K-\$20K.

3.2.2.8.3.5 Schedule

The base version of Secure XENIX is available now. B2 certification is expected to be complete in 2Q89.

3.2.2.8.4 Evaluation

The IBM Secure XENIX option supports the majority of functions required by the TGG. These include adequate performance, low price, and a timely certification schedule. The DDN X.25 suite and TCP/IP protocols are not directly provided by IBM but are available from third party sources. The IBM alternative is lacking in the area of secure network management.

3.2.2.9 SunOS MLS

SUNOS MLS**• ADVANTAGES**

- SECURITY ENHANCED NETWORKING
- PERFORMANCE
- DDN X.25 AND TCP/IP
- TIMELY B1 CERTIFICATION

• DISADVANTAGES

The SunOS MLS operating system is a secure operating system based on 4.3/4.4 BSD and AT&T System V UNIX. SunOS MLS is targeted for a B1 certification expected in 3Q90. The SunOS MLS hardware base consists of the Sun-3, Sun-4, and Sun's TEMPEST workstations.

3.2.2.9.1 Operating System

The SunOS MLS is a multi-level secure operating system targeted for a B1 level certification. SunOS MLS supports access control, multi-level secure window manager, security-enhanced networking, accountability, and secure startup.

For mandatory access control, security labels are associated with all subjects and objects. System administrators or security officers are responsible for the association of labels with subjects and objects. The SunOS MLS security policy enforces no read-up and no write-down as would be expected. The SunOS MLS security policy also prohibits write-up.

The secure window manager is a windowing system and set of utilities that allow users to simultaneously manipulate and display different classifications of data on a single workstation screen. This feature allows users to operate on different classifications of data at the same time without requiring multiple logins and logouts for each classification of data. A secure mailtool utility is available that lets users view various classifications of mail in appropriately labeled windows. With the appropriate privilege, users can also transfer data between windows of varying classifications.

Security-enhanced networking provides the capability for users to exchange labeled information across a network of Sun workstations. A limited security-enhanced capability is supported on networks containing

non-Sun hosts. The capability exists to designate the classification of data received from a non-Sun host. A range of classifications or a single classification may be defined. Sun's network labeling is IP-based.

Special auditing capabilities may be achieved by defining new audit event classes. Audit events may be individually defined for each user to reflect individual audit profiles. This provides the capability of decreasing the amount of auditing information captured for a particular user as that user progresses from a less-trusted to a more-trusted status.

A secure boot PROM is provided that protects the system from accidental or malicious boot attempts.

At present, SunOS MLS does not support the POSIX standard, however plans are to include this feature to the degree possible in the next version of SunOS/MLS.

3.2.2.9.2 Hardware

The Sun 3/150 has been selected as a probable base for the TGG application. These systems run a Motorola 68020 at 16.67 MHz, and are available with 4 to 16 MB of main memory and a variety of disc and tape drives.

A communications co-processor board known as the MCP board is available. This board supports four lines with a maximum aggregate throughput of 500 Kbits/sec.

3.2.2.9.3 TGG Criteria Mapping

The following sections map the capabilities of the SunOS MLS system against the TGG criteria.

3.2.2.9.3.1 Communication Protocols

The MCP board can be used in tandem with Sun's DDN communications software package to support DDN X.25 and TCP/IP. The TCP/IP supports security labeling.

3.2.2.9.3.2 Secure Remote Monitoring and Control

In addition to the standard network capabilities of UNIX, a security enhanced networking capability is available. The security enhanced networking functionality provides for mandatory access control labeling that could lead to a secure remote monitoring and control capability.

3.2.2.9.3.3 Performance

The Sun 3/150 CPU runs at 2 MIPS, and thus our measure predicts a maximum throughput of some 4000 one thousand bit packets per second. As usual, the communications coprocessor board is only capable of supporting a fraction of this load. However the TGG throughput requirements are

attainable with two MCP boards. Given higher line speeds, these two boards could handle on the order of 500 one thousand bit packets per second.

3.2.2.9.3.4 Cost

The base product cost for a suitable base for the TGG should lie in the \$25K - \$45K range, depending on the exact configuration selected, and the degree to which quantity discounts can be arranged.

3.2.2.9.3.5 Schedule

B1 certification of SunOS MLS is expected around 3Q90.

3.2.2.9.4 Evaluation

It appears that there are no disadvantages to the SunOS/MLS approach to the TGG which are not endemic in the trusted operating system approach. The positive side includes cost, support of DDN X.25 and TCP/IP, more than adequate performance, and a timely certification schedule.

3.2.3 *Trusted Operating System Assessment Results*

This section presents the general conclusions regarding the utilization of trusted operating system products functioning as a TGG. Specific recommendations indicating which trusted operating system products are best suited to function as a TGG are also provided.

3.2.3.1 General Conclusions

This section presents findings that apply, as a whole, to trusted operating systems in regards to the TGG criteria. It was considered likely that a subset of the criteria would not be met by the trusted operating systems even before the individual assessments took place. These criteria consisted of gateway operation and remote management protocols and gateway services. To briefly summarize, gateway operation protocols are not supported; remote management is not totally supported, but varying fundamental network functionality to support this capability are available; and the level or granularity of the gateway services in the form of access control and accountability is insufficient.

Communication protocols are generally supported, but not always at the level required for the TGG. Generally, all vendors support X.25, but not all vendors support the DDN X.25 version. Generally, all vendors support TCP/IP. However, many vendors provide TCP/IP packaged with other lower level protocols such as IEEE 802.3 LAN and ETHERNET protocols versus X.25. Also, the support of IP labeling varies widely. In addition, work is needed to provide a gateway version of IP.

Utilizing the approach to determine datagrams/second explained in section 3.2.1.3, Trusted Operating System Performance Calculation, performance does not appear to be a problem. All vendors that have

provided CPU MIPS and I/O processor throughput figures were well above the 200 datagrams/second requirement.

Trusted operating system products generally cost significantly more than commercial gateway products. There were exceptions such as IBM Secure XENIX and Honeywell Secure UNIX. The cost of trusted operating system products was comparable with the one trusted gateway product, the FAC Multinet Gateway.

The certification schedule for all products, with the exception of the BiiN/OS, are well on their way to achieving a B-level certification within the TGG time frame. Certification of the BiiN/OS at a B-level also could possibly be within the TGG time frame. Currently, BiiN plans to obtain B2 level functionality by 4Q90.

3.2.3.2 Product Recommendations

Based on the survey findings of the various trusted operating system products as reflected in Sections 3.2.2.x.3, TGG Criteria Mapping, and Sections 3.2.2.x.4, Evaluation, the following products are most suitable as a baseline for the development of the TGG:

AT&T System V/MLS;
Honeywell Secure UNIX;
IBM Secure XENIX; and
SunOS MLS.

The remaining trusted operating system products were not recommended largely due to the cost factor. Although, the factors of communications protocols and certification schedule did lead to the exclusion of some products from the recommended products list.

3.2.4 Evolution To A TGG

This section provides an example of the development approach that may be utilized for evolving a trusted operating system product to a TGG. In the definition of the development approach, a specific trusted operating system product has been chosen and the subsequent development approach defined in light of that specific product. This approach was chosen, versus a overall trusted operating system product generic approach, to give a better understanding of the actual development specifics. The specific development approach presented for the IBM Secure XENIX operating system is generally applicable to all other trusted operating system products indicated on the recommended list with some expected variations. The following sections detail the efforts necessary to evolve the IBM Secure XENIX operating system to a TGG, including development approach and associated level of effort and cost. The level of effort and cost section primarily focuses on those criteria that must be provided for by additional development efforts. This section

also addresses the cost associated with the evaluation and certification of the TGG.

3.2.4.1 Development Approaches

The total approach for the development of a TGG, utilizing the IBM Secure XENIX product, is presented in the following sections. Consideration is given to all TGG criteria, of a development nature, with a special focus on the critical criteria. The critical criteria, as identified in sections 3.2.1 and 3.2.3, are those TGG criteria that are just simply not supported or that are supported in some form, but not directly utilizable by the TGG. These criteria are the gateway operation protocol, remote management protocol, and gateway services. Communications protocols (i.e., X.25 and TCP/IP) are provided, but in some instances not in the specific configuration required for the TGG or require an additional effort to implement.

The following sections, as delineated by TGG criterion, indicate the implicit features of the IBM Secure XENIX product that satisfy various TGG criteria and how those criteria not provided for by the IBM Secure XENIX product are obtained. In cases where multiple development approaches are applicable for the provision of criteria not supported by the IBM Secure XENIX product, the multiple approaches are presented.

3.2.4.1.1 Protocols

This section presents the methods of acquisition and/or development of the various protocols required for TGG operation. Specifically, these are the DDN X.25, TCP/IP, EGP, and gateway management protocols.

3.2.4.1.1.1 Communication Protocols

DDN X.25 can be provided via the utilization of an optional communications controller board such as Frontier Technologies' Intelligent Advanced Communication Controller (IACC) module. This module provides for the on-board resident implementation of both X.25 (levels II and III) and TCP/IP. The design of the TGG requires that IP be kernel resident. A feature is provided with the IACC module that allows for the disabling of the module resident TCP/IP in support of a kernel-based TCP/IP protocol. In addition to providing a gateway IP product, IP must be reviewed to determine its certifiability and modified appropriately. Also of concern is the review of the X.25 device driver included with the optional communications controller for a determination of its trustworthiness. This is necessary since the device driver was not part of the base evaluated Secure XENIX operating system. Based on that review, modification of the existing device driver or the complete rewrite of the device driver may be necessary to achieve certification.

3.2.4.1.1.2 Gateway Operation (EGP) Protocol

The EGP protocol is not supported by the IBM Secure XENIX operating system. However, the EGP protocol is available as public domain software written in the C programming language and utilizing the features of UNIX. This makes the porting of the EGP protocol possible to the IBM Secure XENIX operating system. Additional efforts will be required to totally assure the correct operation of the EGP protocol and integrate it with the operating system and the TGG application specifics. The integration of the EGP protocol, with Secure XENIX, may become an involved effort due to the security constraints placed on its operation by Secure XENIX. Of extreme importance, is the assurance of the correct operation of the EGP protocol, both from a functional and malicious viewpoint.

3.2.4.1.1.3 Management Protocol

The gateway management protocol is not supported by the IBM Secure XENIX operating system. As with the EGP protocol, public domain software is available that embodies gateway management functionality and will require additional efforts as identified for the EGP protocol. The base gateway management code will require additional functionality in the areas of host pair access rights parameters maintenance, definition of audit parameters (i.e., events to be captured), collection of audit data, and other relevant security actions. A feature that must be provided in conjunction with gateway management is an authentication capability to ensure the trustworthiness of the remote management action source. Authentication may be provided via either COMSEC or COMPUSEC features. The utilization of encryption with digital signatures could provide the necessary source authentication.

The implementation of a future POSIX secure networking standard may provide a base mechanism for the support of the remote management functionality. The current POSIX standard will be evolved to include secure networking services as indicated in Section 3.2.2, Trusted Operating System Survey.

3.2.4.1.2 Gateway Services

This section presents the development approach that leads to the provision of the various gateway services. Specifically, these services are access control and associated flow limitation, auditing, and labeling. All of these services may be provided via the utilization of one of two different development approaches. Labeling is dependent on the particular IP product utilized. Whereas, access control and auditing may be provided via dedicated application programs or via the utilization of Secure XENIX system routines.

3.2.4.1.2.1 Labeling

Labeling is required at the IP level to identify the data coming from a low side network destined for a high side network. This requires processing of IP

security options and the determination of the appropriate label (based on the incoming network interface). This determination can be based on the data within the IP-to-network interface or from system calls which return the label associated with a communication channel. The utilization of system calls to obtain the label or classification of the communication channel provides a more accurate and cleaner solution. Determination of the data label in this fashion ensures the accurate labeling of the data based on the classification of the channel.

3.2.4.1.2.2 Access Control

The provision of the access control and auditing functionality may be provided via a dedicated application program or via the utilization of system routines. A dedicated application program could be utilized that provides access control via the creation, maintenance, and querying of a host access control table. The access control table identifies all hosts including a list of hosts with which an individual host may communicate. This same application program could provide the required auditing functionality and flow limitation. The major advantage associated with this approach is a rather straightforward application-level development effort utilizing application programmers versus specialized programmers such as systems programmers. The major disadvantage involves the need to provide the trusted access checking in an application rather than relying on the existing, trusted operating system. An additional disadvantage is the separate support of an application-level audit trail versus the audit trail provided and protected by Secure XENIX.

Utilization of the implicit features of Secure XENIX provide an alternate approach to the acquisition of access control and associated functionality. The utilization of Secure XENIX features such as account and group definitions and system routines will allow for the provision of the required access control functionality. The following represents an example of the utilization of Secure XENIX features for the provision of access control and its associated functionality:

- definition of accounts based on the identity of the various hosts;
- definition of host groups, all hosts within a group can communicate with each other;
- creation of a rather small, system routine heavy executable that executes within a parent process to perform the following:
 - receives all data;
 - extracts source and destination host parameters (effectively account IDs);
 - changes ownership (chown) of a SETGID executable based on the source host ID;

- updates access control list on the communications device to reflect the destination host's group ID (ACLopen, ACLdel, and ACLadd);
- creates a sub-process (fork); and
- sub-process executes (exec) the SETGID executable mentioned above.
- sub-process SETGID executable is a rather small system routine heavy executable that:
 - reflects the effective GID of the source host required for access arbitration;
 - attempts access to the communications device for transfer of the data to the destination host, the communications device contains an access control list allowing access only to those hosts that are members of the group as indicated by the group ID, which is the group ID of the destination host; and
 - data transfer is either allowed or disallowed and the sub-process terminates.
- the actions are performed repeatedly by the parent process and sub-process are performed for all incoming data.

Both of these processes, or executables that they run, are considered as part of the TGG Trusted Computing Base. Special privileges are required for this approach, however Secure XENIX provides for multiple fine-grained privileges versus a single all powerful privilege. Specifically, the parent process will require the utilization of two privileges for the changing of ownership of the executable that is to be eventually executed by the sub-process and for the alteration of the access control list associated with the communications channel device to reflect the destination host group ID. These required privileges are fine-grained and only utilized for these purposes.

Auditing is implicitly provided during the access arbitration performed by the sub-process. The primary advantage of this approach is the reliance on the operating system for the secure provision of access control and auditing, thereby reducing or eliminating the necessity for application-level code of this type. The major disadvantage of this approach is the requirement for skilled system-level programmers. Also performance may be a problem due to process creation overhead.

3.2.4.2 Development Level of Effort and Cost

This section provides a strawman estimation of the level of effort and associated cost required to evolve the IBM Secure XENIX operating system to a TGG. This estimation is based on prices known for the available hardware

and software base from IBM and Frontier Technologies and on the required development effort as identified in the previous section.

The base Secure XENIX product, which includes a IBM PS/2 Model 80 with an appropriately sized hard disk (i.e., at least 80 MB); the Frontier Technology communications controller(s); and the Secure XENIX operating system, is approximately 15K. This price reflects a recurring cost on a per unit basis. Quantity discounts would further reduce this figure. A one time cost necessary for the acquisition of the TGG is made up of several factors: protocols and gateway services acquisition, specialized services, and certification costs. The protocols and gateway services necessary for the operation of the TGG equate to \$332K to \$367K and \$57K to \$94K respectively. The cost required for specialized services needed during the development effort such as security engineering and management equate to \$216K. The cost required for the certification effort equates to a cost of \$125K bringing the one time cost to a total of \$730K to \$802K. Assuming 50 TGGs are required, the one time cost amount approximately equates to \$15K to \$16K per TGG. Factoring in the reoccurring per unit cost of \$15K, the total acquisition cost per TGG equates to approximately \$30K to \$31K.

The unit cost is exclusive of ongoing life-cycle costs such as maintenance and reflects development costs at a rate of 2.5 times actual salaries to account for the various factors, such as overhead and fee, that figure into the determination of a labor rate. Average salaries are considered to be 35K, 40K, and 45K for an application, communications, and systems programmer respectively. Additional man power costs are factored in for security and management expertise. These costs were factored in at a salary of 50K and 65K respectively. Assuming a one and a half year development period the additional cost for these categories of expertise equate to \$216K. This assumes a half-time participation by these categories of expertise. Approximately a 12 man-month effort will be required for the certification of the TGG utilizing personnel from the security engineering category for a cost of \$125K.

The following sections detail the specifics regarding the necessary level of effort and associated cost for the acquisition of protocols and gateway services.

3.2.4.2.1 Protocols

This section details the level of effort necessary for the acquisition of the various TGG protocols. The individual costs equate to \$18K to \$53K for the DDN X.25 and IP implementation, 100K for the EGP implementation, and \$214K for the remote management implementation efforts for a total protocol implementation cost of \$332K to \$367K.

3.2.4.2.1.1 Communication Protocols

DDN X.25 and IP will be provided as COTS products. Efforts involved concern the review of the X.25 driver and IP software for an indication of

their certifiability. The level of effort for this task may vary dependent on the results of the review, that is modifications may be necessary to the X.25 driver and/or IP software. Modifications would only be necessary to ensure the certifiability of the individual products. Review efforts will require a minimum of one man-month for each product. A systems programmer should be utilized for the review of the driver, whereas a communications programmer should be utilized for the review of the IP software. It may be necessary to expand these efforts by as much as three man-months to account for necessary modifications that may result due to the review. As a result, the X.25 driver task equates to a cost of \$10K to \$28K, whereas the IP task costs from \$8K to \$25K for a total cost of \$18K to \$53K.

3.2.4.2.1.2 Gateway Operation (EGP) Protocol

Efforts involved for the acquisition of the EGP protocol involve the review and integration of the EGP protocol with the other elements of the system (i.e., operating system and TGG application). The level of effort associated with this task equates to approximately a 12 man-month effort at the communications programmer level for a total cost of \$100K.

3.2.4.2.1.3 Management Protocol

Efforts involved for the acquisition of the remote management protocol involve the review and integration of the protocol with the other elements of the system. The effort includes the need to supplement the management protocol with additional functionality such as host access parameters maintenance and various audit capabilities and to incorporate authentication. The level of effort associated with this task equates to approximately a 24 man-month effort at roughly equal partitions for the communications programmer and systems program categories for a total cost of \$214K.

3.2.4.2.2 Gateway Services

This section details the level of effort required for the acquisition of the various gateway services required for the operation of the TGG. The cost of this task, dependent on the development approach utilized, ranges from \$57K to \$94K. IP labeling does not incur any development cost, but may incur a maintenance cost dependent on the mechanism by which the data label is obtained.

As identified in Section 3.2.4.1.2, Gateway Services, there are two methods of providing access control and the associated functions of flow limitation and auditing. The first approach is via the development of a TGG specific application program to provide the access control functionality. The development of the application program can be achieved largely via the utilization of applications programmers. Although, systems programmers should be utilized in a limited role to determine those features of Secure XENIX that may be easily utilized in providing the access control functionality, thereby reducing the amount of applications code required.

The level of effort required to complete this task, utilizing this approach, is approximately 12 man-months utilizing a mix of applications and systems programmers at a ratio of 3:1. Based on this level of effort, the cost to complete this task equates to approximately \$94K.

The second approach involves the utilization of system routines to provide the required access control and associated flow limitation and auditing. A much smaller coding effort is required to provide this functionality via this approach, however the code is much more complicated and requires the efforts of systems programmers. As a result of utilizing this approach, the auditing functionality, as provided by Secure XENIX, is adequate to meet the auditing requirements of the TGG. The level of effort required to complete this task, utilizing this approach, is approximately six man-months at an approximate cost of \$57K. This cost does not include conceptual analysis to further define how operating system primitives would best be used to provide TGG functions.

3.2.4.3 Evolution Summary

As can be seen from the development information presented within these sections, a system based on the IBM Secure XENIX operating system can be easily evolved to a TGG system. The factors that support this statement are the availability of public domain software in support of the various protocols and the variety of ways available by which the gateway services may be acquired. Dependent on the method utilized for the acquisition of gateway services, little or no software development would be required. The evolution of the POSIX standard into the area of secure networking may provide the specific gateway services necessary for the TGG. Worst case would require the utilization of Secure XENIX system routines, with minimal supporting code, to provide a simplified path to the acquisition of the gateway services.

To conclude, UNIX based trusted operating system products as indicated in Section 3.2.3.2, Recommended Products, would provide a base for the TGG that is easily evolvable to a full-scale TGG product. Easily evolvable equating to a relatively minor development and certification effort.

3.3 Gateway Technology

As part of the technology assessment, commercial and government gateway products were surveyed and their applicability to the acquisition or development of a TGG was determined. As noted in our final report for the first phase of this effort, there are no current gateway implementations entirely suited for use as TGGs. However, as we also concluded in that report, the TGG represents only a modest functional departure from commercial and government-sponsored gateway implementations. Therefore, the status of gateway vendors and their products is essential information.

While we do not expect any off-the-shelf product to satisfy the TGG technical requirements, the information exchange between SPARTA and gateway vendors has benefitted all parties. One of the primary focuses of the dialogue with the vendors is to gauge the supplier market for TGGs. In other words, are there many vendors that have products or experience that would be good bases for developing a TGG? Also, if the government were to release an RFP, would there be potential bidders? The interaction with the vendors allows SPARTA to reach conclusions on these issues. At the same time it allows the vendors to gauge the distance between their product lines and plans and the actual TGG requirements.

In addition to reflecting vendors' products and opinions directly, we also provide an independent assessment of the feasibility, cost and issues associated with developing a TGG based upon a commercial gateway product. Gateway vendors' views on this subject were negative, with the exception of Ford Aerospace Corp. On the other hand, our assessment offers evidence for the feasibility of acquiring a TGG via a commercial gateway product.

3.3.1 Vendor Interviews

This phase of the technology assessment involved initiating direct contact with each of the vendors through phone conversations and eventually setting up meetings, frequently at vendor facilities. SPARTA's role in the eventual development and procurement of TGGs was explained as the introductory part of the meeting in almost all cases. Generally, the second part of the meeting was spent apprising the vendors of the TGG requirements and probable program directions. This led naturally to the vendors' presentation of programs and products that they deemed appropriate. SPARTA left this part of the interview up to the discretion of the vendor. At all times, an effort was made to refrain from influencing the vendors in any way. This allowed for the possibility of gathering information on products not originally considered as viable TGG candidates.

The delineation of the SPARTA role at the outset of interviews turned out to be an important part of the process. Frequently vendors were apprehensive as to the motive for gaining detailed information about their product lines. Once assured of the advisory role and the basic preclusion from any implementation role, all were most helpful and even eager to discuss the prospect of TGGs. An area of particular interest to the vendors was the expected market size and demand for trusted gateways both for the specific TGG scenarios and for a wider application. It became apparent that the exchange of information is beneficial for both parties. For SPARTA, it means a more thorough assessment process; for the vendor, information as to what the government might be planning in the future. In only one instance, with Proteon, did a vendor find it necessary to request the execution of a non-disclosure agreement.

Much of the discussion with gateway vendors focused upon product specification and performance. These discussions were guided to a large degree by the checklist found in Figure 3-3. The checklist was a result of the requirements established during the first phase of the contract. It provided a framework for interviewing vendors as well as a basis for comparison between vendors. As was expected, the depth of discussion on each item in the checklist varied with the strengths and weaknesses of each vendor. However, it allowed for the refining of the interview process and was important in the establishment of the ultimate criteria for the gateway technology assessment.

GATEWAY FUNCTIONS

- ACCESS CONTROL
 - HOST PAIR BASIS
 - UPPER LEVEL APPLICATION BASIS
- IPSO PROCESSING
- FLOW LIMITATION
- DDN INTERFACES
- AUDIT CAPABILITIES

GATEWAY PERFORMANCE

- PACKETS/DATAGRAMS PER SECOND
- OTHER MEASURES

COST PER UNIT

- DEPENDENCE UPON CONFIGURATION, NETS SUPPORTED, ETC.

MONITORING AND CONTROL

- PROPRIETARY APPROACHES
- SUPPORT FOR EMERGING STANDARDS (SNMP, SGMP, HEMS)

CERTIFICATION

- PLANS TO CERTIFY ANY PRODUCTS
- OBSTACLES SEEN IN CERTIFICATION PROCESS
- TARGET CERTIFICATION LEVELS
- SOFTWARE ARCHITECTURE OF CERTIFICATION
- HARDWARE ARCHITECTURE OF CERTIFICATION

PROTOCOL SUITE SUPPORTED

- ISO IP VS. DoD IP
- EGP
- GGP
- ETC.

ESTIMATED TIME AND EFFORT TO DEVELOP TGG

Figure 3-3 Preliminary Checklist for Gateway Vendor Discussion

3.3.2 Criteria

In the final report for phase 1 of this effort, some initial conclusions were drawn as to what the more important issues in TGG acquisition and development might be, namely certification and performance. At the outset of this second phase, the issues expected to bear most heavily upon the timely, cost efficient development and production of a TGG were once again examined and the most significant ones were used as criteria against which to measure each vendor. As outlined and described below, some of those issues that we originally ranked as critical factors (e.g., performance and certain kinds of technology) are regarded by many vendors as simple matters of programming. Other issues that raised concern (e.g., certification and evaluation) continue to be primary factors requiring serious attention.

The criteria discussed below were used to identify definite areas of strength or preference among potential gateway vendors. Given an equal rating in other areas, a difference in one of these areas would indicate that that vendor is better suited to build a TGG. We have not at this time assigned weights to the criteria in order to compare vendors with differing strength profiles.

3.3.2.1 Schedule

A primary consideration in this area was whether the vendor could support an aggressive schedule, taking into account such factors as certification status of products and the vendor's knowledge of the technology and marketplace. For a clearer presentation, the certification issue is separated from the schedule issue in the vendor profiles that follow. These factors would affect not only the types and numbers of responses to a TGG procurement effort, but would impact the overall development effort timeframe.

3.3.2.1.1 Certification

The conclusion reached in the final report for phase 1 of this effort is that the TGG poses a significant certification effort. Each vendor's plans (if any) to certify products and their respective target certification levels were examined. If there were no current plans for a certification effort of an appropriate product, the vendor's willingness to undertake and support a certification process as well as their understanding of such a process was determined. At this point in the technology assessment, certification remains the largest obstacle to the fielding of the TGG.

Some vendors desire the ability to claim a certified product among their product line, but many are concerned about the time and dollar investment required for such an effort. For one vendor in particular, Proteon, it could be a matter of committing corporate talent to a project that would handicap their potential for growth in other areas.

3.3.2.1.2 Technology

An attempt has been made to determine whether the prospective vendors have the requisite knowledge and understanding of DoD protocols, security and certification issues. Only minor differences among vendor capabilities are expected, given their status as competing gateway developers. However, the assessment found commercial off-the-shelf (COTS) gateways having superior performance to the FAC Multinet Gateway, despite the MNG's use of multiple processors.

3.3.2.1.3 Marketplace

The concern in this area is whether the vendors are familiar with the DoD contracting cycle as well as the certification and evaluation process. Lack of familiarity could easily result in program delays that could otherwise be avoided. This can often be determined by examining their current customer base and product lines to find correlation between that and the kind of project that the TGG procurement is likely to resemble.

3.3.2.2 Protocols

Whether the vendor, and the product in question, actually supports DoD protocols and the degree to which they support them is a category in and of itself. Also part of this question, is whether the vendor is flexible. For example, a vendor's product may not totally support DoD protocols, but the fact that they have significant knowledge and express an interest in doing so would be a factor. The ideal TGG gateway vendor would provide complete support of DoD protocols in addition to the intent to make the transition/migration to ISO protocols.

3.3.2.3 Network Management

Given the absence of established network management standards, the vendor should be familiar with different techniques for network management approaches and be capable of rapid implementation once such standards are defined and specified. An area of great concern is the need for strong authentication for network management control messages.

3.3.2.4 Special Services

Access control and several other special services are required to be supported by the TGG. The application of this criteria means measuring what special services a vendor may offer, if any, as well as determining how close they may be to implementing these services if they are not in existence. Besides access control, audit trail generation and flow limitation capabilities are necessary and even more important is that the design and implementation of these services must support the eventual certification of a product as they are the services that enforce the TGG-specific security policy.

3.3.2.5 Performance

Although predicted TGG performance is important as an assessment factor, strict comparison of the obtained estimates is not advocated. COTS vendors report performance using non-secure operating systems, while FAC reports performance using a secure operating system. There is also a significant variation in the definition of performance measures and in the techniques used to measure or estimate them. Analysis performed in phase 1 of this contract determined that a performance number on the order of 200 datagrams per second is called for from the TGG. It should be noted that this number is a suggested estimate rather than a firm requirement.

3.3.2.6 Cost

The vendors were not always capable of providing accurate figures for future implementations that might take into account the required added TGG functionality. However, given the knowledge gathered from all of the vendors surveyed to date, it was possible to estimate unit costs within reasonable bounds. For the basis of comparison, purchasing a single TGG for approximately \$50K seems to be the current high bracket. More detailed discussion is required about the tradeoffs between unit costs and development costs in the case of adding functionality (e.g., incorporating a trusted operating system in a gateway product).

While the cost factor is bound to be a very important parameter in the procurement process, strict comparisons of the cost data reported below would be misleading as costs of COTS vendor products are for volume-produced units and do not reflect the cost of a TGG that the same vendor might produce. Only the cost of the FAC Multinet Gateway represents a cost for a product including both the functionality and certification features required. The aggregate cost data do provide a spectrum between \$12,000 and \$50,000.

3.3.3 Gateway Vendor Survey

The following narratives describe what level each vendor currently reaches under each of the criteria. Each criterion is discussed for each interviewed vendor. A summary is also provided for each vendor.

3.3.3.1 Proteon

**PROTEON
p 4x00 SERIES GATEWAY**

• ADVANTAGES

- EXISTING HIGH PERFORMANCE PRODUCT FOR DoD INTERNET**
- LARGE INSTALLED PRODUCT BASE**
- POTENTIAL FOR MEETING TGG FUNCTIONAL REQUIREMENTS**

• DISADVANTAGES

- CURRENT PRODUCT PROBABLY NOT CERTIFIABLE**
- VENDOR UNABLE TO SEEK CERTIFICATION ALONE**

Proteon's product line includes gateways, internetwork bridges as well as interface devices for their proprietary LANs. Their installed product base numbers in the thousands.

3.3.3.1.1 Certification

Proteon is very reluctant to support unassisted certification of their gateway software. Proteon is unfamiliar with the process, but not the concepts of certification. They expressed an interest in undertaking a hypothetical TGG development with a system-integrator partner able to provide support services beyond the basic development and production phases. This is quite a feasible approach for Proteon, because they are planning to team with a large system house on a project with requirements similar to those of the TGG.

Proteon's software is based upon the public domain CMOS operating system. They have enhanced it for speed and additional functionality. CMOS is inherently simple and tightly structured¹. We believe that a B1 - certifiable operating system could be developed from CMOS for on the order of \$300-500K. Therefore, the certifiability of Proteon products, in terms of cost and schedule, poses a moderate rather than a very high risk. Sections 3.3.4.3 and 3.3.4.4 develop cost estimates and list issues associated with this potential development.

3.3.3.1.2 Schedule

¹Berglass, G.R. "CMOS, A Portable Operating System in C," MTR-84-W-00071, APR 1984

Proteon could support an aggressive TGG development schedule with proper assistance. They envision themselves primarily as suppliers of COTS products, but are willing to undertake custom development work. As a commercial gateway vendor, Proteon is planning product improvements along the lines of TGG security requirements; however, they are unprepared to develop a certified product without assistance.

3.3.3.1.3 Protocols Support

Proteon supports the DoD protocol suite and participates in DoD committees such as the Internet Engineering Task Force (IETF). They have both the current capability to support DoD protocols now and the flexibility to evolve to the ISO set of protocols. Proteon also supports several LAN protocols, indicating their expertise in this area.

3.3.3.1.4 Network Management

Proteon is at the leading edge in developing network management services. They currently provide network management services including SGMP and, soon, SNMP; they participate in network management standards committees; and they have developed and demonstrated OverVIEW, a PC-based network management system.

3.3.3.1.5 Special Services

Proteon provides only limited special services, such as IP-pair based access control. However, they consistently upgrade their software releases and can respond easily to added requirements and specifications.

3.3.3.1.6 Performance

Proteon has consistently strived for performance via software upgrades and movement to increasingly capable hardware bases. Their planned port to a VME RISC-based CPU leads them to expect processing rates of 10,000 datagrams or packets per second.

3.3.3.1.7 Cost

Prices for Proteon's gateways vary with the features included (e.g., Ethernet boards, X.25 interfaces, etc.). Their average shipped value per unit is \$18,000.

3.3.3.1.8 Conclusions

Although Proteon is a strong gateway developer, significant concerns exist about their ability to be the TGG developer. They are strong because they have a well-known product with a large installed base and because they could easily develop gateway software to meet the TGG functional requirements. However, their current product would be difficult to certify in view of its several years of evolution via incremental performance improvements and feature additions. In addition, Proteon has strong reservations about responsibility for meeting the non-technical requirements associated with the DoD marketplace: certification and logistic support. They would prefer to

have a large partner take those responsibilities. Their current marketing plans may present them with such a partnership opportunity.

3.3.3.2 CISCO

CISCO MODEL AGS-1E1D GATEWAY

- ADVANTAGES

- EXISTING HIGH PERFORMANCE PRODUCT FOR DoD INTERNET
- INSTALLED BASE
- POTENTIAL FOR MEETING TGG FUNCTIONAL REQUIREMENTS

- DISADVANTAGE

- VENDOR HAS LITTLE INTEREST IN DEVELOPMENT CONTRACT
- VENDOR HAS LITTLE DESIRE FOR PRODUCT CERTIFICATION

CISCO's product line is COTS and is currently limited to gateways that can have interfaces with high-speed (e.g., Ethernets) and X.25 networks. They are well known and have installed many system for DoD customers.

3.3.3.2.1 Certification

CISCO is not eager to undertake the certification process. Conversations with CISCO uncovered the fact that CISCO wishes to sell COTS equipment rather than to be a development contractor. As the latter activity is clearly required for TGG certification, a plan for certification would be a significant problem using the CISCO gateway as the TGG basis. However, CISCO is currently teaming with Honeywell to develop a TEMPEST gateway, showing their willingness to team with others to develop products to meet security requirements. Furthermore, CISCO has identified private sector security requirements similar to DoD requirements.

CISCO's software and operating system is proprietary. They noted that it "does contain a process control model." However, it is not possible to know whether it would lend itself to development of a certifiable operating system. Therefore, CISCO's operating system does not offer any means of risk reduction or schedule enhancement.

3.3.3.2.2 Schedule

CISCO is planning product improvements very similar to TGG security requirements, and they are competent software engineers. They expressed confidence that they could easily develop a gateway with the required TGG

features. Therefore, using a CISCO gateway as a TGG basis would probably support very rapid technical development of TGG features.

3.3.3.2.3 Protocols Support

CISCO products run the DoD protocol suite, including DDN X.25 and CISCO participates in DoD-oriented committees such as the IETF. They have both the current capability to support DoD protocols as well as the flexibility and plans to evolve to ISO protocols.

3.3.3.2.4 Network Management

CISCO currently provides limited network management services, and they participate in network management standards committees. Based on CISCO's basic competence and their participation in standards development, they could develop expanded network management as required for the TGG.

3.3.3.2.5 Special Services

CISCO provides limited granularity access control via IP source addresses; transit of datagrams may be refused on the basis of IP address pairs. CISCO is capable of other special service developments but is not eager to undertake them under government contract.

3.3.3.2.6 Performance

CISCO products use the Motorola 680x0 series processors together with a tailored operating system; their current performance approaches 1,000 datagrams per second between attached Ethernets.

3.3.3.2.7 Cost

A CISCO gateway for interconnecting DDN X.25 with an Ethernet (the AGS-1E1D) is priced at \$12,200.

3.3.3.2.8 Conclusions

CISCO is a strong gateway developer, but they have little desire to be a development contractor. They have shown their ability to develop and deliver high-performance gateways by their current installed base. Their development skills could easily be applied to meeting the TGG functional requirements. However, their expressed desire is to sell COTS gateways, perhaps to a second party who could support certification and logistics requirements.

3.3.3.3 Ford Aerospace Corporation (FAC) Multinet Gateway (MNG)

**FORD AEROSPACE CORPORATION
MULTINET GATEWAY**

• **ADVANTAGES**

- COULD SUPPORT ENCRYPTED DISNET TRAFFIC ACROSS MILNET
- WOULD SUPPORT TGG FUNCTIONAL REQUIREMENTS
- IS CERTIFIED GATEWAY PRODUCT

• **DISADVANTAGE**

- COST - \$50K/UNIT
- LOW PERFORMANCE - 100 DATAGRAMS/SEC.
- STILL IN DEVELOPMENT
- NO REMOTE MANAGEMENT YET

Ford Aerospace Corporation (FAC) has developed the Multinet Gateway (MNG) over the past 6-7 years under contract to RADC and supporting other DoD customers. Several advanced development models have been placed with DoD customers, but no large scale production runs have been made. The Multinet Gateway program includes several versions of the product: the original advance development model currently undergoing evaluation in conjunction with the Global Decision Support System; a Model 1 planned for production as a commercial product and which does not include cryptography; an enhancement of the Model 1 including cryptography under the Commercial COMSEC Endorsement Program; and other possible versions. In conducting this assessment, the Model 1 is chosen as the most likely base for a TGG.

3.3.3.3.1 Certification

The MNG provides interconnection among multiple networks operating at multiple security levels. It has been submitted for A1 certification and is currently under evaluation. This experience places the product and vendor significantly ahead of other vendors with respect to development of a certified product. Therefore, the cost and schedule would present low risks.

3.3.3.3.2 Schedule

FAC can potentially apply its development experience and its planned expansion of the MNG product line to development of a TGG. The technical development schedule could be quite aggressive, given the short distance to be spanned between the MNG and the TGG. However, FAC has little experience with large scale production and developing this capability could pose some risk to an aggressive schedule.

3.3.3.3.3 Protocol Support

The MNG currently supports DoD gateway protocol standards, including the Exterior Gateway Protocol (EGP). They have no current plans to evolve to ISO protocol standards. If the market dictates, FAC could be persuaded to develop ISO protocols.

3.3.3.3.4 Network Management

The MNG contains "stubs" for routines to perform secure remote management, but no features are currently included. These stubs are also intended to support remote management only by another MNG. FAC has bid to work on a secure monitoring center capability using a MNG as a part of the Enhanced Multinet Gateway Program. If this work is awarded to FAC, the resulting experience could be applied to TGG management requirements. The use of evolving management protocol standards including secure management would represent a new effort for FAC.

3.3.3.3.5 Special Services

The MNG offers a subset of special services required by the TGG. These include access controls based upon TCP header information, and "guard" functions for supporting the sending of UNCLASSIFIED data by an UNCLASSIFIED process to a SECRET process. The access control capability defined for the TGG has previously been implemented as an application running on the MNG as a part of previous support for DCA. The labeling of IP datagrams as to the trustworthiness of the source (in accordance with TGG requirements) could be developed easily from existing MNG functions. The current MNG supports IPSO labeling in order to label UNCLASSIFIED datagrams received without labels. This capability can be modified to add the extended IPSO label required for the TGG. Based on these existing features, the MNG can readily support special gateway functions needed to meet TGG requirements.

3.3.3.3.6 Performance

The MNG supports on the order of 100 datagrams per second aggregate throughput using all of its ports simultaneously. This performance is significantly less than what is offered by COTS gateway vendors.

3.3.3.3.7 Price

A single MNG unit would have an estimated cost of \$50,000.

3.3.3.3.8 Conclusions

The MNG represents a product very close to the TGG in terms of its functions and in terms of certification requirements. In addition, its MNG functions could support exchange of DISNET encrypted traffic across the UNCLASSIFIED MILNET segment. There are disadvantages associated with the MNG as well: its performance is only around 100 datagrams per second; it is still in development, without benefit of experience with an installed base; and there is no current way of remotely managing a MNG. Its cost is significantly higher than that of some of the commercial gateways.

3.3.3.4 Bolt, Beranek & Newman (BBN) Butterfly Mailbridge

BOLT BERANEK & NEWMAN BUTTERFLY MAILBRIDGE

- **ADVANTAGES**
 - EXISTING ACCESS CONTROL FUNCTIONS
 - DDN COMPLIANT GW AND MC OPERATION
 - PATH TO AUTHENTICATED MANAGEMENT
 - GOOD PERFORMANCE
- **DISADVANTAGES**
 - DIFFICULT PATH TO CERTIFICATION
 - HIGH END OF TARGET COST RANGE

BBN has developed several generations of network and internetwork packet switching systems used within the Defense Data Network. The Butterfly gateway was developed from their generic (Butterfly) multiprocessor architecture. The Butterfly machine has been demonstrated and proposed for a number of other applications, such as scientific and artificial intelligence computing. The Butterfly Mailbridge Gateway has been developed to provide controlled intersegment operation within the DDN with functional requirements very similar to the TGG but without any requirements for certification. The Mailbridge serves as the basis for this assessment.

3.3.3.4.1 Certification

BBN has no plans at this time for developing a trusted computing base in conjunction with their Butterfly gateway and its Chrysalis operating system. Chrysalis supports the management of multiple processes within a shared memory model and facilitates memory accesses by processes.

Consequently, it would be very difficult to base a trusted kernel on Chrysalis. We discussed with BBN the possibility of porting their gateway software to an alternative operating system. BBN felt that the close coupling of the gateway software with the Chrysalis design would result in severe performance penalties (estimated at an order of magnitude) in the event of such a port. These factors imply significant cost and schedule risks associated with certifying the BBN Butterfly gateway.

3.3.3.4.2 Schedule

An aggressive development schedule of a TGG based upon the Butterfly gateway could be pursued, if certification issues are not considered. This is because straight-forward extensions of Butterfly features and functions can meet TGG technical requirements. BBN is well known for its work in developing network software and hardware.

3.3.3.4.3 Protocols Support

The BBN Butterfly gateway supports the current DoD gateway protocols and is in use within laboratory environments to support next generation gateway protocols. A conversion to ISO protocols is pending as well.

3.3.3.4.4 Network Management

Butterfly gateways currently serve as DDN inter-segment gateways between the MILNET and the ARPANET. They support all of the current DDN monitoring and Gateway-to-Gateway protocols. BBN's experience in developing and managing elements in the DDN suggests that TGG network management requirements would not pose significant technical risks.

3.3.3.4.5 Special Services

BBN may add authentication to the Butterfly features to assure that current network management traffic is safe. Also, the Butterfly gateway can perform the access control services needed by a TGG based upon rules selected at compile time. There is no current support for dynamic access control.

3.3.3.4.6 Performance

The Butterfly gateway is capable of performance on the order of 3,000 datagrams per second in a 16-CPU machine.

3.3.3.4.7 Price

The cost of a single unit Butterfly Mailbridge Gateway has been estimated at \$50,000.

3.3.3.4.8 Conclusions

The BBN Mailbridge is a high performance gateway with functions that could be easily extended to meet TGG technical requirements. The current access control functions are close to the TGG requirements, and its protocols for Gateway-to-Gateway and Monitoring Center interaction are DDN compliant. The Butterfly offers an easy path to authenticated management,

and it offers strong performance. However, the prospects for developing a certified gateway from the Butterfly are remote due to the nature of the current operating system and shared memory. Further consideration of the BBN gateway needs to concentrate on assessing options for providing the requisite trusted computing base. In addition, the Butterfly Gateway cost is significantly higher than the cost of a standard commercial gateway.

3.3.3.5 Other Gateway Vendors

Attempts were made to establish contact with other gateway vendors with product lines oriented to the DoD or Federal market. Several phone calls were placed to Communication Machinery Corporation (CMC) and to 3COM (who bought Bridge Communications). The difficulty in obtaining responses from these vendors suggests that their interests lie elsewhere, similar to the cases of Proteon and CISCO. These are both volume-oriented vendors with competitively priced products in the range of \$10,000 - \$20,000 per unit.

Both of these vendors' products were noted in the Phase I report of this contract. The products can interconnect Ethernet LANs and Wide Area Networks using DoD protocols. Our interest in Bridge was heightened by reports that it has developed encryption features for its gateways. However, Bridge was bought by 3COM and the current status of the encryption-based product is unknown.

3.3.4 Summary

The survey of gateway vendors uncovered both strengths and weaknesses in this potential TGG supplier segment. The strengths lie in the availability of commercial high-performance gateways; the weakness is in the difficulty that would be faced in certifying commercial gateway products. These ideas are expanded below, leading to the conclusion that a partnership may offer the best prospects for acquiring TGGs.

We also present an independent assessment of the costs and issues associated with developing a TGG from a commercial product. Our assessment identifies steps for enhancing an existing secure gateway and for enhancing and certifying a commercial gateway product. These estimates are offered only as a consequence of some knowledge and understanding about the operating system structure of a commercial gateway product-- the Proteon series which use the CMOS operating system. Similarly, the discussion of issues in providing a trusted operating system depends explicitly on our awareness of the CMOS operating system's structure.

3.3.4.1 Scarcity of Cross-qualified Vendors

Despite the growth of networking technology during the past two decades, there are few vendors who have a combination of gateway development, COMSEC, and COMPUSEC expertise needed for TGG

development. A simple reason for this may be the difference between vendors' major markets, which demand high speed at a competitive price, and the DoD classified subscriber community which requires certifiably secure systems. It is difficult for any vendor to develop a strong presence in both of these markets.

Vendors who sell larger numbers of gateways (e.g., Proteon and CISCO) make profits based upon their volume. They perceive a risk in devoting the corporate talent and resources that would be necessary for developing the TGG; such a commitment might preclude them from future higher-volume, high-profit opportunities. Understandably, these vendors have a definite interest in participation in a TGG development, but they would prefer a partner capable of shouldering the necessary burdens of certification and support.

There is certainly a risk in not considering experienced gateway vendors for the TGG acquisition. DoD needs the expertise associated with developing gateways and probably cannot afford to have a TGG vendor learn about gateway development and operation through on-the-job training. The risk would be the amount of learning a non-gateway vendor would require to attain competence in gateway development.

3.3.4.2 Need for Partnerships and Roles

Given the need for an experienced gateway developer, but the unsuitability of commercial gateway code for a certified TGG, a partnership offers the best prospects for acquiring TGGs based on commercial gateway products. The partnership must include an experienced gateway developer for the obvious reasons: gateways are high-performance real-time systems, and the TGG will require combinations of porting and redevelopment. The partnership must also include a trusted operating system supplier. Much effort and time can be saved by using a trusted operating system as a basis for certifying the TGG. Ideally, the trusted operating system vendor can also lead the TGG certification process. The partnership should also include a system integrator who can be responsive to the operational needs of DISNET and other DoD network subscribers. The integrator's skills should include training and maintenance support.

3.3.4.3 TGG Development Estimates

The following sections provide estimates for the costs associated with developing a TGG based on a gateway product. These cost assume a well defined specification for the TGG and do not include substantial resources for system engineering efforts to design a TGG. A more detailed discussion of the nature of the steps required to produce a TGG is presented in section 3.2.4. Our estimates cover two phases: enhancing an existing secure gateway, and enhancing and certifying a gateway product. Both phases are necessary to

develop a TGG from a commercial gateway product; therefore the cost of the former must be added to the cost of the latter in any final analysis.

3.3.4.3.1 Cost Estimate for Enhancing an Existing Secure Gateway

An existing secure gateway product, such as the FAC Multinet Gateway runs as a set of processes using a Trusted Computing Base. To develop it into a TGG requires that the datagram labeling and access control functions be added to the existing software. Our cost estimates are based upon the following work element breakdown:

1 Man Mo.	Modify the gateway's forwarding function to invoke the labeling and access control functions;
6 Man Mos.	Implement the Access Control function, including provisions for flow limitations;
2 Man Mos.	Implement the labeling function (i.e., via the IP Security Option);
3 Man Mos.	Augment the existing management functions to control the access control and labeling functions
12 Man Mos.	Augment the existing management functions to include authentication of management commands
6 Man Mos.	Documentation of the Augmented Gateway functions to support certification of the TGG configuration
<hr/>	
30 Man Mos.	TOTAL for developing enhancements to an existing secure gateway.

This development is a non-recurring cost. For purposes of this analysis, a man month is approximately equal to \$10,000 of development cost. This is sufficiently accurate for 1989 to support the comparisons among acquisition paths. Our cost calculation will include the \$50,000 baseline recurring per-unit cost associated with the FAC Multinet Gateway, described in Section 3.3.3.3. The costs for quantities of 20 TGGs, including development costs, would be \$65K per unit and would be \$55K per unit for quantities of 50.

3.3.4.3.2 Cost Estimate for Enhancing and Certifying a Gateway Product

An existing commercial (non-secure) gateway product will require enhancements to its gateway software as described above, and it will require restructuring of its underlying operating system to provide a basis for certification at the B1 level. Our cost estimates for the operating system restructuring are described below. There is a wide choice of real-time operating systems capable of supporting gateway functions. Within this variety, there may be operating systems that lend themselves easily to

restructuring and ones that present greater difficulty. (Sibert, et al.¹ discuss the case of building a secure system based upon UNIX.) We provide a range estimate for activities devoted to developing a secure system based upon a native operating system. Our cost estimates are based upon the following work element breakdown:

3 Man Mos.	Define a security architecture and assess the Operating System and its Kernel functions for their suitability for restructuring into a given security architecture.
6 - 24 Man Mos.	Modify and enhance the Operating System Kernel functions to meet the security architecture requirements
6 - 24 Man Mos.	Modify the existing gateway software to handle revisions in the operating system, including possible restructuring of the process set;
12 Man Mos.	Documentation of the Operating System Structure and of the augmented gateway functions
30 Man Mos.	Steps defined in section 3.3.4.3.1
-----	-----
57 - 93 Man Mos.	TOTAL for activities to develop a TGG based upon an existing non-secure gateway.

This development is a non-recurring cost. For purposes of this analysis, a man month is approximately equal to \$10,000 of development cost. This is sufficiently accurate for 1989 to support the comparisons among acquisition paths. Our cost calculation will include the \$18,000 baseline recurring per-unit cost associated with the Proteon Gateway, described in section 3.3.3.1. The costs for quantities of 20 TGGs, including development costs, would be \$60K per unit and would be \$35K per unit for quantities of 50.

In addition to these costs is the cost of certification. The steps necessary to support certification, particularly documentation, are included in the above estimates. The actual certification is not included. Estimates for that depend on whether the certification is performed by the NCSC, is contracted out, or is performed by the developer.

3.3.4.4 Gateway Operating System Security

One of the critical areas for developing a TGG based on an existing commercial gateway product is the gateway operating system. Most of the

¹Sibert, W. O., Traxler, H. M., Wagner, G. M., Downs, D. Elliot, K. B., and Glass, J. J., "UNIX and B2: Are They Compatible?"

commercial gateways examined as a part of this assessment use a minimal real time operating system selected to support the gateway performance. The gateway application is closely coupled with the underlying operating system in order to minimize overhead. In order to meet TGG requirements, the gateway operating system must be modified to meet the requirements for a B1 operating system. This step is included in the previous section under the need to modify and enhance the operating system.

3.3.4.4.1 Rationale for a Trusted Gateway Operating System

One of the critical areas for developing a TGG based on an existing gateway product is the issue of a trusted operating system. The trusted operating system provides the foundation for insuring that the TGG performs as intended. A trusted operating system provides a number of services for the TGG.

First, it can protect the gateway software, both applications and operating system, from tampering. By enforcing controlled access, the executing portions of the TGG can be segregated from other untrusted applications and from user processes. This is an essential part of the overall configuration control approach which includes the configuration control for the operating system.

The operating system also provides assurance that labeling and separation are preserved. Separation involves high and low side network data; monitoring, control, and management information; gateway operational data; and any user process information. This separation capability also provides assurance that the access control functions are invoked. By assigning labels to different functions, the operating system can require the invocation of a trusted process to move data from one side to another. Further, the operating system can support the IPSO based labeling.

Associated with the operating system is support for auditing. Depending on how TGG functions are integrated, the operating system audit capabilities can also serve to collect and protect TGG audit data. The testing performed as a part of the operating system evaluation and certification can serve as a basis for the TGG security testing.

Finally, if there is a provision for local or remote user interfaces, a trusted operating system supports user identification, authentication, access control, and separation.

3.3.4.4.2 Issues in Providing a Trusted Gateway Operating System

The previous section has highlighted the importance of incorporating a trusted operating system as a part of the TGG. In addition, the cost estimates have allocated resources for modifying the operating system in commercial gateways to meet B1 criteria. While certifying an existing system for security requirements is unpopular and uneconomical, circumstances sometimes leave no alternatives. The following discussion of issues addresses what

must be done when circumstances dictate that an existing system must be certified to meet a B1 assurance level. The issues are based upon requirements imposed for B1 certification; in several cases an immediate closure can be suggested, and in others the issue is left open. This issue analysis suggests that certifying an existing system is feasible and possibly economically viable. For each of the major B1-specific requirements, we note a major issue associated with enhancing a gateway operating system to meet the requirement. We also note whether we expect the issue to remain open or whether it should be quickly resolvable in an enhancement effort aimed at evolving a commercial gateway to a TGG. An example of such an enhancement which might address these issues would be the modification of the CMOS operating system in the Proteon gateways.

1. Requirement for Discretionary Access Control among named subjects and objects

OPEN ISSUE, BUT EVENTUAL RESOLUTION EXPECTED

It is more difficult to implement this requirement in packet-switching equipment, because objects are created and disposed of at high rates. Names must be used and re-used, and access control list activity can become oppressive. Object naming should be handled via inheritance, such as channel of entry into the system. This is an open issue, but it can be resolved by simple concensus, possibly using approaches employed in the MultiNet gateway.

2. Requirement for assurance against subject access to data for which it has no authorization as a consequence of object re-use

ISSUE SUBJECT TO QUICK RESOLUTION

This issue can be closed by provision of a TCB service for object allocation and de-allocation. There is a performance penalty for this solution, however.

3. Requirement for sensitivity labeling of objects, subjects, capabilities

OPEN ISSUE REQUIRING DESIGN STUDY

This requirement is met by use of the IP Security Option and by labeling capabilities of devices and channels. The header checksum method must be demonstrated to assure that the sensitivity label is not subject to undetected errors. Subjects (i.e., processes) must possess similar labels.

OPEN ISSUE REQUIRING CONTRACTOR-GOVERNMENT COORDINATION

Policies must be established for human-readable output generated by the TGG (regardless of the avenue of development).

4. Requirement for mandatory access control, including both hierarchical and category sets

OPEN ISSUE REQUIRING DESIGN STUDY

An efficient mechanism, as part of the TCB, is required for mediating each access, performing the access control logic check and permitting the access only when the logic succeeds. Additionally, this requirement can be addressed via system design changes that provide strictly-labeled working space for the labeled subjects. Access confined to the labeled working space need not be mediated.

5. Requirement to authenticate users**OPEN ISSUE REQUIRING CONTRACTOR-GOVERNMENT COORDINATION**

Categories of users must be determined, regardless of the avenue of TGG development.

OPEN ISSUE REQUIRING DESIGN STUDY

Mechanisms are needed to unambiguously and uniquely identify any remote entities with claims to access control privileges or to system management privileges.

6. Requirement to provide audit capabilities**OPEN ISSUE, BUT EVENTUAL RESOLUTION EXPECTED**

Auditable events must be defined, and auditing functions must be developed and added.

7. Requirement for a separate domain of operation for the TCB**ISSUE RESOLVABLE THROUGH AVAILABE HARDWARE FEATURES AND DESIGN STUDY**

The Motorola 68020 central processing unit provides 'supervisor' and 'user' processing modes. The Motorola 68851 Paged Memory Management Unit provides a hierarchical address space along with enforcement of process access rights. Together these hardware features can provide a certifiable basis for meeting this requirement.

8. Requirement for process isolation**ISSUE RESOLVABLE THROUGH AVAILABE HARDWARE FEATURES AND DESIGN STUDY**

Mechanisms for interprocess communication must be developed within the TCB. They must enforce the information security policies for reading and writing.

9. Requirement for hardware validation of TCB integrity**OPEN ISSUE, BUT EVENTUAL RESOLUTION EXPECTED**

Self-test methods must be developed to provide assurance that TCB-critical data and code have no undetected errors.

In conclusion, while the enhancement of an existing operating system to meet B1 requirements will involve an appreciable effort, such a modification should be achievable. The emphasis on features at the B1 level, as contrasted with assurance and architecture at the B2 level, make a B1 gateway operating system a practical choice. If the TGG required a B2 operating, such a retrofit would probably not be cost effective.

3.4 Technology Assessment Summary

The technology assessment considered a range of options for a base for acquiring a TGG. The options included trusted operating systems for a wide range of hardware bases as well as a number of gateway products. With the exception of the Multinet Gateway program, no option represented a close match with the TGG functional and security requirements. Within the operating system and gateway categories a great deal of similarity was seen in terms of the effort required to enhance products to become TGGs.

3.4.1 Assessment Overview

In the operating system category, as expected, all the options require the porting or development of the gateway functionality for that implementation option including the specialized TGG services. In general, the trusted operating system vendors were interested in expanding into the network market and showed significant interest in the TGG direction. Secure remote management is a problem for all vendors, as well as for all the gateway vendors, although some had approaches for management of homogeneous systems. The secure management issue is likely to be resolved by incorporating evolving standards for management protocols including strong authentication. Within the operating system category, the selection of leading candidates is largely based upon an assessment of the ease of supporting gateway functionality and consideration of the cost and performance factors.

Within the gateway category, the major issue was one of certification. Most gateway products are implemented around real time operating systems where both the application and operating system emphasize performance. The gateway vendors tend to be major market oriented and are reticent to pursue unilaterally a special case. Certification is viewed as a major effort requiring significant resources. As a consequence, the gateway vendors expressed greater interest in a cooperative effort in which a teammate could provide logistics and computer security support for a TGG program. Despite these reservations, having a trusted gateway product in their inventory was viewed as desirable. In terms of functionality, many of the gateway vendors were implementing variations on many of the TGG special services. The performance of all the commercial gateway products was well beyond that required for a TGG within the DDN. The selection of leading candidates among gateways is based upon an assessment of which is most conducive to

and is willing to proceed towards certification with some consideration also given to cost and functionality.

The Multinet Gateway (MNG) is the only example of a current trusted gateway product. This is an ongoing development effort to provide secure gateway services. Functionally, it either directly meets the TGG requirements or requires minor changes, however there remain a number of questions concerning the applicability of the MNG. From a performance standpoint it is significantly slower than any other alternative and is below the target performance range for the TGG on the DDN. The MNG is also still a custom development item and there remain reservations concerning the transition to commercial production of the MNG and its eventual cost. Within these bounds, as the only existing trusted gateway, the MNG serves as a basis against which other candidates need to be measured.

3.4.2 Conclusions

Sections 3.2 and 3.3 provided detailed assessments of alternatives for procuring a trusted guard gateway. The options considered were almost all capable of being modified and enhanced to meet the functional requirements described in section 2. In some cases, this represented minor changes and straightforward development while in others substantial effort would be needed. In considering an overall assessment, the options can be categorized in terms of the ability to meet functional requirements, the resulting performance, the risk associated with meeting the TGG requirements including certification, and the unit cost.

A rating has been assigned in each category for each alternative considered in the technology assessment. These ratings are based on combinations of the objective characterizations determined from interactions with vendors and on subjective assessments developed during the survey. In each category, a rating of 1 is best. Combining the ratings into an overall rating for each alternative is based on the relative importance of each category. We have not attempted to assign such weightings.

Cost is a rating based on the estimated unit cost of a TGG based on that technology option. The cost is based on a likely hardware configuration and does not reflect special developmental or other non-recurring costs. A rating of 1 indicates a cost of \$0-30K; a rating of 2 indicates a cost of \$30-60K; and a rating of 3 indicates a cost of greater than \$60K.

Performance is a rating based on the throughput values for gateway products and on our estimates for operating systems. Our approach for operating system throughput calculation is described in section 3.2.1.3. These estimates, and the ratings assigned, are meant only to be order of magnitude ranges and not precise characterizations. A rating of 1 indicates throughput greater than 1000 datagrams/second; a rating of 2 indicates throughput of between 200 and 1000 datagrams/second; and a rating of 3 indicates a throughput of less than 200 datagrams/second and is below the performance

level established for a TGG. The performance ranges are based on the basic hardware and software for the product. Additional limitations may be imposed as a consequence of the I/O coprocessor performance. At this time, only SUNOS MLS and BiiN have I/O coprocessors identified which will support the processor performance. For the TGG, those alternatives with ratings of 1 or 2 meet the performance needs of the environments described in section 2.

Functionality is an assessment of the magnitude of effort associated with meeting functional requirements for the TGG and with producing a well behaved gateway. This category reflects the basic distinction between gateway and operating system products and indicates the expected ease with which a trusted operating system could support gateway functionality. A rating of 1 indicates minor modifications or small additions; a rating of 2 indicates a matter of routine development; and a rating of 3 indicates a need for a major effort. These ratings, along with risk, tend to reflect the likely magnitude of the non-recurring, or developmental, costs associated with each alternative.

Risk is the most subjective category. This category reflects the concerns associated with certification along with concerns about the impact of incorporating gateway functions. Risk tends to be higher of operating systems which have not been widely used for networking applications and also for hardware bases which involve distributed processors. Lower risk ratings also tend to indicate a greater likelihood of meeting an aggressive schedule. A rating of 1 indicates low risk; a rating of 2 indicates modest risk; and a rating of 3 indicates the highest risk.

VENDOR	UNIT COST	PERFORMANCE	FUNCTIONALITY	RISK
AT&T SYSTEM V/MLS	2	1	2	2
BIIN/OS	3	1	3	3
GEMINI GEMSOS/PCAT	2	2	3	3
GOULD UTX/32S	3	1	2	2
HONEYWELL STOP	3	2	2	3
HONEYWELL SECURE UNIX	2	1	2	2
IBM SECURE XENIX	1	1	2	2
SUNOS MLS	2	1	2	1
PROTEON	1	1	1	3
CISCO	1	2	1	3
BBN BUTTERFLY MB	2	2	1	3
FAC MULTINET GATEWAY	2	3	1	1

Figure 3-4 TGG Technology Alternative Ratings

Based on these assessments, the leading candidates for a TGG are IBM's Secure XENIX, SUNOS MLS, Proteon, and the FAC Multinet Gateway. A selection among these alternatives involves the relative importance of risk, performance, and cost. These alternatives also do not address possible cooperative efforts which could produce low cost, low risk, high performance alternatives.

4.0 APPLICATION RELAY APPROACH

4.1 *Introduction*

In light of the recent events associated with the internet worm (or virus) incident, the desire has emerged to provide a finer granularity of access control than that currently included in the TGG. In order to realize this desire and to provide a more robust firewall, an alternative architecture for the TGG is being developed. Rather than an IP gateway with enhanced access control services, this alternative is a Trusted Application Relay (TAR).




Figure 4-1. Trusted Application Relay Architecture

The TAR is a dual-homed host that terminates all end-to-end protocols, including applications such as file transfer or mail. A relay function connects two independent protocol stacks and performs the security related functions. A separate relay function would be needed to connect each application type. The number of such relay functions required would depend upon the number of applications that are authorized. The basic services are the same as

those in the TGG, but are now provided at the application level. This permits labeling at the application level and the enforcement of access control rules based on application specific information. In addition, the TAR provides high confidence authentication as a result of the termination of end-to-end protocols.

Although the TAR is envisioned as connecting similar protocol stacks, it could also provide application translation services by connecting two different stacks. This might support OSI transition services such as FTAM to FTP relays or X.400 to SMTP relays. This potential coupling of missions could provide cost advantages. In addition, the TAR could operate in conjunction with a layer 3 relay, a TGG as defined in section 2. This TGG would provide the IP connectivity for a subset of hosts on a host address basis and would provide the only virtual terminal support (i.e., TELNET or VTP). Such a TGG operating in this type of combination could either be a separate piece of hardware or could be integrated into the TAR as a cutout in the protocol stack.

The following sections provide a description of the characteristics and services associated with the TAR. This definition has been refined as a part of the overall TGG analysis effort and this section considers implementation choices for a TAR based upon the candidate products identified in section 3.

4.2 Application Relay Security Services

The motivation for the consideration of a TAR is to provide finer granularity security services than found in the TGG with a higher degree of confidence that those services are performed as intended. The nature of these services derive from the information available as a consequence of processing the application specific information. The TAR has the capability of enforcing limitations on the services invoked within an application. The following sections describe specific security functions that a TAR could perform.

4.2.1 Authentication

The degree of authentication that can be provided through the TAR with existing applications is still limited. True user authentication within mail or file transfer applications requires enhancements to those protocols such as those being defined for privacy enhanced mail¹. The first class of improved authentication that the TAR offers compared with the TGG is the ability to definitely identify the application being performed. In the TGG case, the TGG must rely on proper use of assigned TCP port numbers to identify the application. For example, there is no way the TGG can recognize conspiring hosts using the assigned mail port number to use a virtual terminal protocol. In the TAR case, because the application is terminated, the nature of the application is definitely established.

¹ "RFC 1040: Privacy Enhancements for Internet Electronic Mail," Jan 1988

The second class of authentication is weaker. Where applicable, the TAR could validate the correspondence of address information at different levels. Addresses exist explicitly or implicitly at least at the network, internetwork, and application levels. For traffic not passing through a relay, the TAR could validate the IP and mail addresses. In addition, for local network traffic, the TAR could validate the network and internetwork addresses. In general, the utility of this type of validation is extremely limited and is probably not useful in a TAR.

4.2.2 Application Specific Access Control

The primary motivation for preferring the TAR architecture is to provide access control and filtering at the application level. Rather than simply limiting which applications may be used between host pairs, the TAR can check and limit application level header fields and contents, application commands, and transfer modes. Further, these rules can be represented in a flexible and extendible method to allow an arbitrary set of policies to be enforced. This flexibility would allow a rapid reaction to a discovered vulnerability in a protocol. For example, in the recent worm incident, the TAR could have been set to prevent addresses containing a pipe character. This could have curtailed intersegment contamination.

Possible rules that might be enforced on mail traffic (e.g., on SMTP) besides address field limits include the checking of HELO information; restricting use of VRFY, EXPN, SEND, and related commands; and the limiting of responses from high side hosts.

For file transfer, the TAR can more clearly identify the direction of the data transfer and the nature of the transaction. In addition to providing greater confidence in the rules defined for the TGG, the TAR could (for FTP, for example) limit binary or image transfer; restrict the use of APPE, MSND, ALLO, RNFR, DELE, LIST, and related commands; and possibly restrict specific filenames such as common UNIX commands. These rules might be enforced on a direction (high side or low side) or host address basis.

Additional applications, particularly transaction based directory operations are likely to need to be supported.

The more extensive set of rules involved in these scenarios will require a larger and probably more dynamic database. An efficient means for representing and checking the rules along with the caching of decisions would need to be developed.

4.2.3 Labeling

The TGG is required to support labeling of datagrams originating in the low side as being of "suspicious origin". This label would have to be processed by the destination host's IP handler and somehow relayed to the end user for useful filtering. Alternatively, the destination host would have to unilaterally either pass or discard all such datagrams. In the TAR, labels

could also be placed within mail messages. By using the encapsulation approaches for messages, the TAR could include an ASCII warning in the message that the recipients could then directly take into account.

4.2.4 Guard Functions

One of the services identified as desirable for the TGG is an upgrade/downgrade function. While the downgrading of text messages or any other data specific processing is still considered beyond the state of the art for the general case, the termination of end-to-end protocols does allow for some limited upgrade capability. If the TAR is at least a B2 certified system, then it could pass mail traffic from an UNCLASSIFIED host to a SECRET host with no acknowledgement or other feedback to the UNCLASSIFIED host. The only impact of this extension of services would be to increase the required level of COMPUSEC certification for the TAR.

4.3 *Implications of the TAR Architecture*

The security features described above serve as the motivation for a TAR. Those features must be balanced against the mechanisms in a TAR needed to implement them. This section presents some of the impacts on a TAR that would distinguish it from the simpler TGG.

As a dual homed host terminating all end-to-end protocols, the TAR acts as a store and forward relay for all mail and file transfer processing. This requires significantly increased processing and storage. The storage requirements are at least those needed to provide for interim reassembly and storage of mail or files prior to relaying. If the TAR buffers information when a high side host is down and awaits its return, then the storage requirements can become very large. A compromise could be developed in which the low side host receives an acknowledgement only if a connection to the high side host succeeds. This could limit, but not eliminate, the storage requirements. In any case, the TAR requires mass storage support not needed by the TGG. Even though it is a non-real time service (as are mail and file transfer, in general) the TAR requires significantly more processing power and main memory in order to keep up with all the connections active at a given time.

Perhaps the most important characteristic of the TAR is that it is distinctly NOT transparent. The non-transparency is apparent both in terms of added delays in application response and in a change in the way in which applications are used. For mail applications, this is a minor issue since the mail system is intended to accommodate store and forward, non-real time relays. File transfer, however, has no such capability to support intermediate relays. In such a case changes to the protocol or to the way in which the protocol is used would be required. File transfer users would have to adapt to a more batch oriented view of file transfer where a request was issued with no confirmation or a delayed confirmation. In addition to the application

specific non-transparency, the TAR would effectively hide the network(s) on the high side, which has both advantages and disadvantages.

4.4 Technology Assessment

The technology assessment described in section 3 presents an assortment of products that might be used for a TAR. The criteria upon which the evaluation is based assume a TGG application. The criteria for a TAR will be somewhat different. Despite those differences, the catalog of products provides a basis for assessing options for a TAR as well. The criteria differences are discussed below.

The TAR role as a dual-homed host effectively rules out the gateway products. The gateway products are included for the TGG assessment since gateway functionality is an important criterion. Gateway functions are no longer required for the TAR. Trusted operating system products provide both the host functions and the certification that are required and will consequently form the list for consideration for a TAR.

Changed criteria involve protocols and performance. Support for gateway protocols is no longer required. Management could be handled through either authenticated management protocols or through secure remote login. A full protocol suite including all authorized applications must be supported. This includes at least IP, TCP, SMTP, and FTP (and/or their ISO equivalents). From a performance standpoint the TAR base must be a more powerful machine than that needed to support TGG performance requirements. It must have increased primary and secondary mass storage to support application level protocol processing. Performance needs to be measured in terms of raw processing power, I/O capabilities, context switching rate, and multi-"user" support.

The impacts of these differences on rating technology options is to opt towards the more powerful trusted operating system hardware configurations. The unit cost for a TAR should be on the order of \$40-60K based on choices such as SUNOS MLS or AT&T System V/MLS. BiiN also is a potential candidate for a TAR if the certification issue (BiiN's plans for a B-level system) is resolved. The following estimate is for the effort required to develop a TAR

36 Man Mos.	Develop application relay functions
12 Man Mos.	Modify protocols to support application relay functions
12 Man Mos.	Provide authentication for remote management or remote login capability
18 Man Mos.	Provide documentation for application relay functions needed for operation and to support certification
12 Man Mos.	Extend remote management capability to manage specific access control functions and rules associated with TAR
90 Man Mos.	TOTAL for developing a TAR from a trusted operating system product

Including development costs in the TAR price, this equates to approximately \$100K per TAR for a quantity of 20 and \$70K per TAR for a quantity of 50. The cost difference between these figures and those of the TGG are primarily based on the increased base cost of the hardware.

4.5 TAR Summary

This approach is presented for consideration, but not as a definite recommendation. At this time there are a number of options for providing protection for hosts engaged in intersegment communication. There are similarly a number of options for implementing whichever architecture is deemed appropriate. The intent of this report is to present the information with which alternatives can be evaluated and selected. As with open issues described elsewhere in the report, further exploration and resolution of these questions is expected through continuing analysis and discussions with DCA and other DoD representatives.